# AWS Storage Gateway

## User Guide

## API Version 2013-06-30

# AWS Storage Gateway: User Guide

Copyright © 2015 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

# Table of Contents


What Is AWS Storage Gateway? ..... 1
- Are You a First-Time AWS Storage Gateway User? ..... 2
- How AWS Storage Gateway Works ..... 2
  - Volume Gateway Architecture (Gateway-Cached and Gateway-Stored) ..... 2
  - Gateway-VTL Architecture ..... 5
- Available AWS Regions ..... 7
- Pricing ..... 8
- Accessing AWS Storage Gateway ..... 8

Getting Started ..... 9
- Recommended Reading ..... 9
- Plan Your Gateway Deployment ..... 10
- Requirements ..... 11
  - Supported Hypervisors and Host Requirements ..... 11
  - Hardware Requirements ..... 11
  - Supported iSCSI Initiators ..... 12
  - Supported Backup Software (Gateway-VTL Only) ..... 12
  - Storage Requirements ..... 12
  - Network Requirements ..... 12
- Step 1: Sign Up ..... 14
- Step 2: Deploy a VM and Activate the Gateway ..... 14
  - On-Premises Gateway ..... 15
  - EC2 Gateway ..... 57

Volume Gateways ..... 66
- Related Topics ..... 66
- Setting Up Your Volume Gateway ..... 67
  - Step 1: Before You Begin ..... 67
  - Step 2: Configure Local Storage and Alarms ..... 67
  - Step 3: Create Volumes ..... 74
  - Step 4: Connect Your Volumes to Your Windows Client ..... 79
  - Step 5: Test the Setup ..... 86
  - Step 6: Clean Up After the Example Exercise ..... 94
  - Where Do I Go from Here? ..... 94
- Managing Your Activated Gateway ..... 96
  - Managing Storage Volumes ..... 97
  - Connecting to Storage Volumes ..... 111
  - Performing Maintenance Tasks ..... 125
  - Monitoring Your Gateway ..... 185
  - Logging API Calls by Using CloudTrail ..... 200
  - Working with Snapshots ..... 201
  - Troubleshooting ..... 226

Gateway-VTL ..... 240
- Setting Up Your Gateway-VTL ..... 240
  - Step 1: Before You Begin ..... 241
  - Step 2: Configure Local Storage ..... 241
  - Step 3: Create Virtual Tapes ..... 243
  - Step 4: Connect VTL Devices to Your Windows Client ..... 245
  - Step 5: Test the Setup ..... 249
  - Step 6: Clean Up After the Example Exercise ..... 275
  - Where Do I Go from Here? ..... 276
- Managing Your Activated Gateway-VTL ..... 276
  - Managing Tapes ..... 277
  - Connecting to VTL Devices ..... 288
  - Performing Gateway-VTL Maintenance ..... 302
  - Monitoring Your Gateway-VTL ..... 351
  - Logging API Calls by Using CloudTrail ..... 361

Troubleshooting ..... 363
Access Control ..... 373
Example Policies ..... 374
API Reference ..... 379
Required Request Headers ..... 379
Signing Requests ..... 381
Example Signature Calculation ..... 381
Error Responses ..... 383
Exceptions ..... 383
Operation Error Codes ..... 384
Error Responses ..... 402
Operations ..... 403
Appendices ..... 404
Appendix A: Components in Your vSphere Environment ..... 404
Appendix B: Configuring a VMware ESXi Host ..... 406
Appendix C: Components in Your Hyper-V Environment ..... 410
Appendix D: Configuring a Hyper-V Host ..... 411
Setting Up and Configuring a Microsoft Hyper-V Host ..... 412
Troubleshooting Your Microsoft Hyper-V Setup ..... 417
Appendix E: AWS Storage Gateway Resources ..... 421
Appendix F: AWS Storage Gateway Limits ..... 421
Appendix G: Launching and Activating an Amazon EC2 AMI into a Non-Default VPC ..... 422
Appendix H: VTL Devices ..... 423
Selecting a Medium Changer After Gateway Activation ..... 424
Updating the Device Driver for Your Medium Changer ..... 425
Appendix I: Symantec Backup Exec 2012 ..... 426
Configuring Storage in Backup Exec 2012 ..... 426
Importing a Tape in Backup Exec 2012 ..... 427
Writing Data to a Tape in Backup Exec 2012 ..... 428
Archiving a Tape Using Backup Exec 2012 ..... 428
Restoring Data from a Tape Archived in Backup Exec 2012 ..... 429
Disabling a Tape Drive in Backup Exec 2012 ..... 429
Document History ..... 430

# What Is AWS Storage Gateway?

**Topics**


- Are You a First-Time AWS Storage Gateway User? (p. 2)
- How AWS Storage Gateway Works (Architecture) (p. 2)
- Available AWS Regions (p. 7)
- AWS Storage Gateway Pricing (p. 8)
- Accessing AWS Storage Gateway (p. 8)


Welcome to the *AWS Storage Gateway User Guide*. AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to provide seamless integration with data security features between your on-premises IT environment and the Amazon Web Services (AWS) storage infrastructure. You can use the service to store data in the AWS cloud for scalable and cost-effective storage that helps maintain data security. AWS Storage Gateway offers both volume-based and tape-based storage solutions:

- **Volume gateways** – Volume gateways provide cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers. The gateway supports the following volume configurations:
  - **Gateway-cached volumes** – You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Gateway-cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.
  - **Gateway-stored volumes** – If you need low-latency access to your entire data set, you can configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive off-site backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2.
- **Gateway–virtual tape library (VTL)** – You can cost-effectively and durably archive backup data in Amazon Glacier. Gateway-VTL provides a virtual tape infrastructure that scales seamlessly with your business needs and eliminates the operational burden of provisioning, scaling, and maintaining a physical tape infrastructure.

You may choose to run AWS Storage Gateway either on-premises, as a virtual machine (VM) appliance, or in AWS, as an EC2 instance. You deploy your gateway on an EC2 instance to provision iSCSI storage volumes in AWS. Gateways hosted on EC2 instances can be used for disaster recovery, data mirroring, and providing storage for applications hosted on Amazon EC2.

For an architectural overview, see How AWS Storage Gateway Works (Architecture) (p. 2).

AWS Storage Gateway enables a wide range of use cases. For more information, go to the AWS Storage Gateway detail page.

# Are You a First-Time AWS Storage Gateway User?

The preceding section summarizes the storage offerings provided by the AWS Storage Gateway. For a detailed architectural overview of these offerings, see How AWS Storage Gateway Works (Architecture) (p. 2).

This documentation provides a Getting Started section that covers setup information common to all gateways and also gateway-specific setup sections. The Getting Started section shows you how to deploy and activate any gateway. The gateway-specific sections then discuss performing further setup for different gateway types: volume gateways (that is, gateway-cached and gateway-stored volumes) and gateway-VTL. Together, the Getting Started section and the gateway-specific sections show you how to get started, set up a functioning gateway, and manage your gateway.

- Getting Started with AWS Storage Gateway (p. 9), for all gateways

- Volume Gateways (Gateway-Cached and Gateway-Stored) (p. 66)
  - Setting Up Your Volume Gateway (Gateway-Cached and Gateway-Stored) (p. 67)
  - Managing Your Activated Gateway (Gateway-Cached and Gateway-Stored) (p. 96)

- Gateway–Virtual Tape Library (Gateway-VTL) (p. 240)
  - Setting Up Your Gateway-VTL (p. 240)
  - Managing Your Activated Gateway-VTL (p. 276)

The instructions in this guide primarily show the gateway operations by using the AWS Management Console. If you want to perform these operations programmatically, see the AWS Storage Gateway API Reference for information about the supported operations.

# How AWS Storage Gateway Works (Architecture)

This section provides an architectural overview of the available AWS Storage Gateway solutions.

**Topics**

- Volume Gateway Architecture (Gateway-Cached and Gateway-Stored) (p. 2)
- Gateway–Virtual Tape Library (Gateway-VTL) Architecture (p. 5)

## Volume Gateway Architecture (Gateway-Cached and Gateway-Stored)

AWS Storage Gateway offers two configurations for volume gateways. In both cases, all data is securely stored in AWS. The primary difference is how much data is stored on-premises.

**Topics**

- Gateway-Cached Volume Architecture (p. 3)

- Gateway-Stored Volume Architecture (p. 4)

# Gateway-Cached Volume Architecture

Gateway-cached volumes let you use Amazon Simple Storage Service (Amazon S3) as your primary data storage while retaining frequently accessed data locally in your storage gateway. Gateway-cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their frequently accessed data. You can create storage volumes up to 32 TiB in size and attach to them as iSCSI devices from your on-premises application servers. Your gateway stores data that you write to these volumes in Amazon S3 and retains recently read data in your on-premises storage gateway's cache and upload buffer storage.

Gateway-cached volumes can range from 1 GiB to 32 TiB in size and must be rounded to the nearest GiB. Each gateway configured for gateway-cached volumes can support up to 20 volumes and a total volume storage of 150 TiB.

In the gateway-cached volume solution, AWS Storage Gateway stores all your on-premises application data in a storage volume in Amazon S3.

The following diagram provides an overview of the AWS Storage Gateway-cached volume deployment.

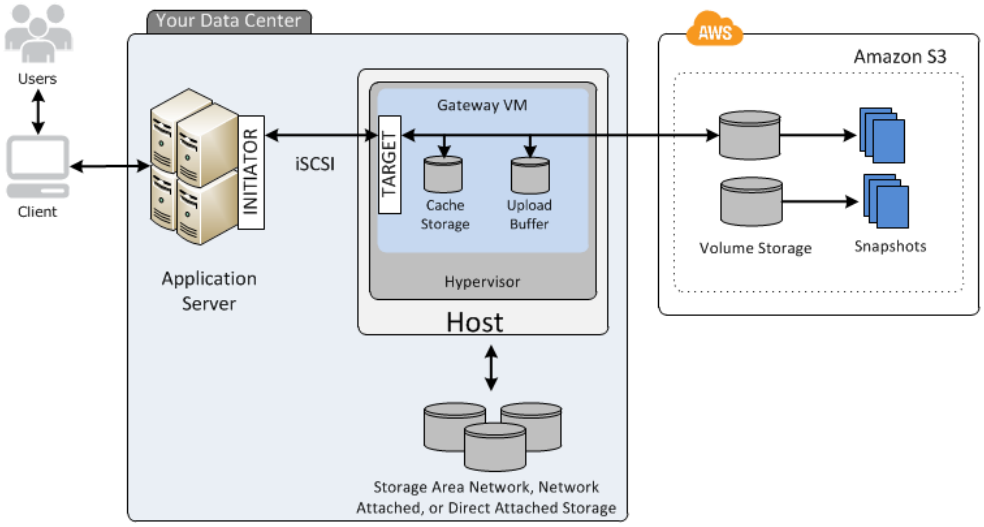


After you've installed the AWS Storage Gateway software appliance—the virtual machine (VM)—on a host in your data center and activated it, you can use the AWS Management Console to provision storage volumes backed by Amazon S3. You can also provision storage volumes programmatically using the AWS Storage Gateway API or the AWS SDK libraries. You then mount these storage volumes to your on-premises application servers as iSCSI devices.

You also allocate disks on-premises for the VM. These on-premises disks serve the following purposes:

- **Disks for use by the gateway as cache storage** – As your applications write data to the storage volumes in AWS, the gateway initially stores the data on the on-premises disks referred to as cache storage before uploading the data to Amazon S3. The cache storage acts as the on-premises durable store for data that is waiting to upload to Amazon S3 from the upload buffer.

  The cache storage also lets the gateway store your application's recently accessed data on-premises for low-latency access. If your application requests data, the gateway first checks the cache storage for the data before checking Amazon S3.

You can use the following guidelines to determine the amount of disk space to allocate for cache storage. Generally, you should allocate at least 20 percent of your existing file store size as cache storage. Cache storage should also be larger than the upload buffer. This latter guideline helps ensure cache storage is large enough to persistently hold all data in the upload buffer that has not yet been uploaded to Amazon S3.

- **Disks for use by the gateway as the upload buffer** – To prepare for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as an *upload buffer.* Your gateway uploads this buffer data over an encrypted Secure Sockets Layer (SSL) connection to AWS, where it is stored encrypted in Amazon S3.

You can take incremental backups, called *snapshots*, of your storage volumes in Amazon S3. These point-in-time snapshots are also stored in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. You can initiate snapshots on a scheduled or one-time basis. When you delete a snapshot, only the data not needed for any other snapshots is removed.

You can restore an Amazon EBS snapshot to a gateway storage volume if you need to recover a backup of your data. Alternatively, for snapshots up to 16 TiB in size, you can use the snapshot as a starting point for a new Amazon EBS volume. You can then attach this new Amazon EBS volume to an Amazon EC2 instance.

You can restore an Amazon EBS snapshot to a gateway storage volume if you need to recover a backup of your data. Alternatively, for snapshots up to 16 TiB in size, you can use the snapshot as a starting point for a new Amazon EBS volume. You can then attach this new Amazon EBS volume to an Amazon EC2 instance.

All gateway-cached volume data and snapshot data is stored in Amazon S3 encrypted at rest using server-side encryption (SSE). However, you cannot access this data with the Amazon S3 API or other tools such as the Amazon S3 console.

## Gateway-Stored Volume Architecture

Gateway-stored volumes let you store your primary data locally, while asynchronously backing up that data to AWS. Gateway-stored volumes provide your on-premises applications with low-latency access to their entire datasets, while providing durable, off-site backups. You can create storage volumes and mount them as iSCSI devices from your on-premises application servers. Data written to your gateway-stored volumes is stored on your on-premises storage hardware. This data is asynchronously backed up to Amazon Simple Storage Service (Amazon S3) in the form of Amazon Elastic Block Store (Amazon EBS) snapshots.

Gateway-stored volumes can range from 1 GiB to 16 TiB in size and must be rounded to the nearest GiB. Each gateway configured for gateway-stored volumes can support up to 12 volumes and a total volume storage of 192 TiB.

In the gateway-stored volume solution, you maintain your volume storage on-premises in your data center. That is, you store all your application data on your on-premises storage hardware. Then, using features that help maintain data security, the gateway uploads data to the AWS cloud for cost-effective backup and rapid disaster recovery. This solution is ideal if you want to keep data locally on-premises, because you need to have low-latency access to all your data, and also to maintain backups in AWS.

The following diagram provides an overview of the AWS Storage Gateway–stored volume deployment.

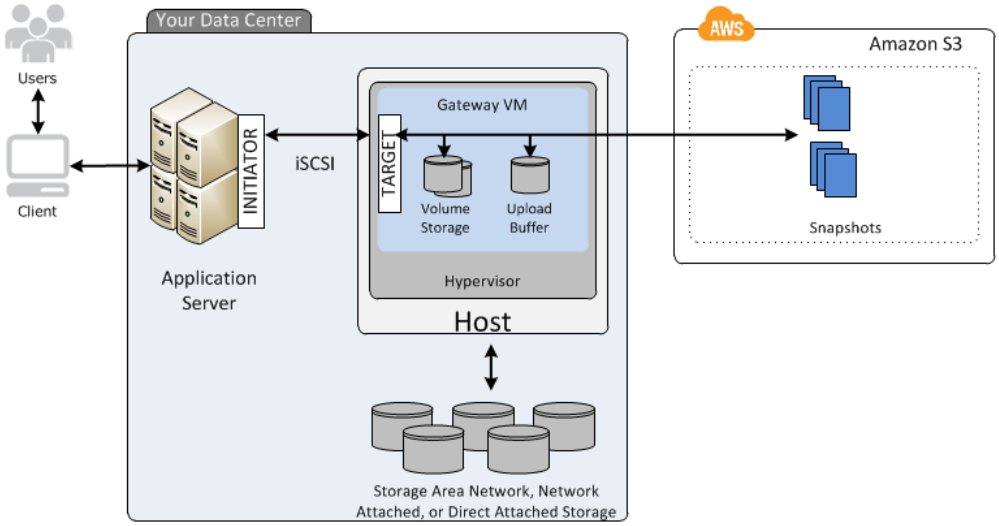


After you've installed the AWS Storage Gateway software appliance—the virtual machine (VM)—on a host in your data center and activated it, you can create gateway *storage volumes* and map them to on-premises direct-attached storage (DAS) or storage area network (SAN) disks. You can start with either new disks or disks already holding data. You can then mount these storage volumes to your on-premises application servers as iSCSI devices. As your on-premises applications write data to and read data from a gateway's storage volume, this data is stored and retrieved from the volume's assigned disk.

To prepare data for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as an *upload buffer*. You can use on-premises DAS or SAN disks for working storage. Your gateway uploads data from the upload buffer over an encrypted Secure Sockets Layer (SSL) connection to the AWS Storage Gateway service running in the AWS cloud. The service then stores the data encrypted in Amazon S3.

You can take incremental backups, called *snapshots*, of your storage volumes. The gateway stores these snapshots in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. You can initiate snapshots on a scheduled or one-time basis. When you delete a snapshot, only the data not needed for any other snapshot is removed.

You can restore an Amazon EBS snapshot to an on-premises gateway storage volume if you need to recover a backup of your data. You can also use the snapshot as a starting point for a new Amazon EBS volume, which you can then attach to an Amazon Elastic Compute Cloud (Amazon EC2) instance.

# Gateway–Virtual Tape Library (Gateway-VTL) Architecture

Gateway-VTL offers a durable, cost-effective solution to archive your data in the AWS cloud. The VTL interface it provides lets you leverage your existing tape-based backup application infrastructure to store data on virtual tape cartridges that you create on your gateway-VTL. Each gateway-VTL is preconfigured with a media changer and tape drives, which are available to your existing client backup applications as iSCSI devices. You add tape cartridges as you need to archive your data.

The following diagram provides an overview of the AWS Storage Gateway-VTL deployment.

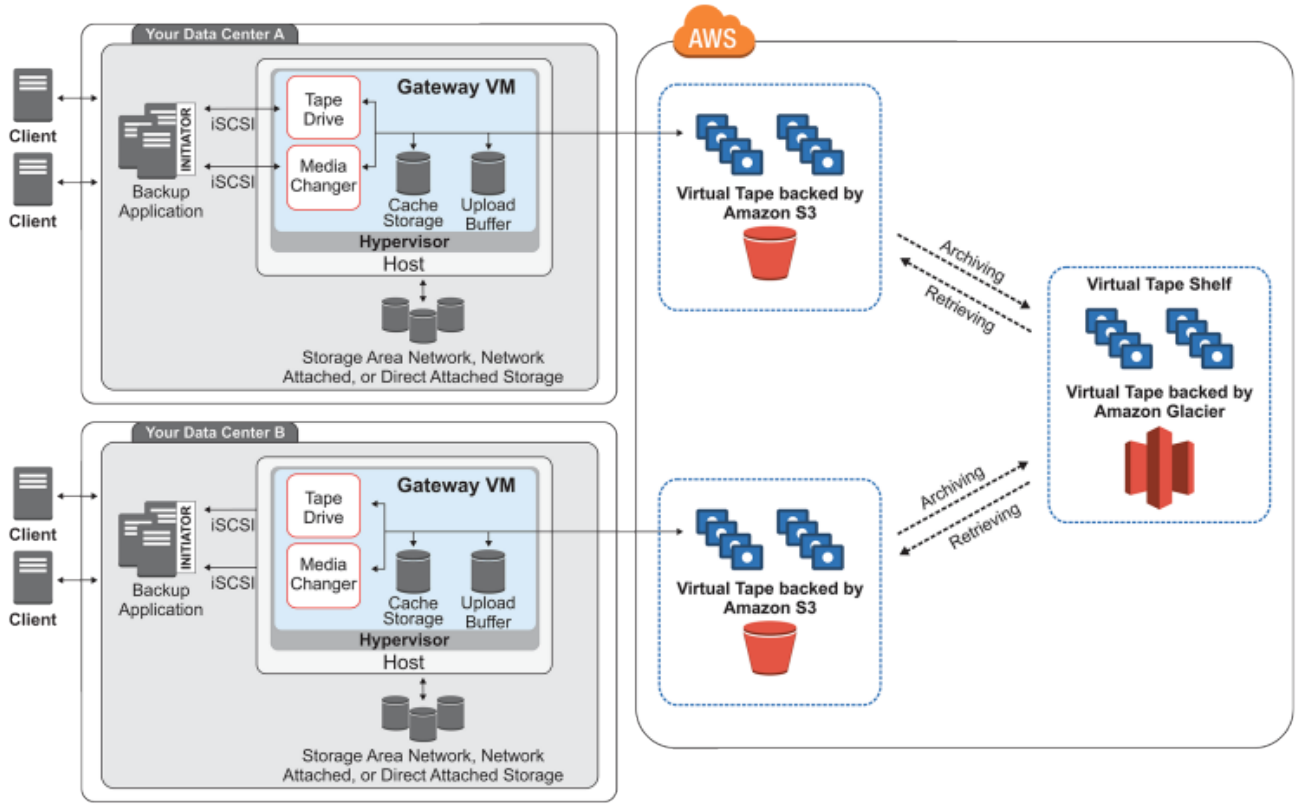


sim

The diagram identifies the following gateway-VTL components:

- **Virtual tape** – Virtual tape is analogous to a physical tape cartridge. However, virtual tape data is stored in the AWS cloud. Like physical tapes, virtual tapes can be blank or can have data written on them. You can create virtual tapes either by using the AWS Storage Gateway console or programmatically by using the AWS Storage Gateway API. Each gateway can contain up to 1500 tapes or up to 150 TiB of total tape data at a time. The size of each virtual tape, which you can configure when you create the tape, is between 100 GiB and 2.5 TiB.
- **Virtual tape library (VTL)** – A VTL is analogous to a physical tape library available on-premises with robotic arms and tape drives, including the collection of virtual tapes stored within the library. Each gateway-VTL comes with one VTL.

  The virtual tapes that you create appear in your gateway's VTL. Tapes in the VTL are backed up by Amazon S3. As your backup software writes data to the gateway, the gateway stores data locally and then asynchronously uploads it to virtual tapes in your VTL—that is, Amazon Simple Storage Service (Amazon S3).
  - **Tape drive** – A VTL tape drive is analogous to a physical tape drive that can perform I/O and seek operations on a tape. Each VTL comes with a set of 10 tape drives, which are available to your backup application as iSCSI devices.
  - **Media changer** – A VTL media changer is analogous to a robot that moves tapes around in a physical tape library's storage slots and tape drives. Each VTL comes with one media changer, which is available to your backup application as an iSCSI device.
- **Virtual tape shelf (VTS)** – A VTS is analogous to an off-site tape holding facility. You can archive tapes from your gateway's VTL to the VTS and, if needed, retrieve tapes from the VTS back to your gateway's VTL.
  - **Archiving tapes** – When your backup software ejects a tape, your gateway moves the tape to the VTS for long-term storage. The VTS is located in the AWS region in which you activated the gateway.

Tapes in the VTS are stored in Amazon Glacier, an extremely low-cost storage service for data archiving and backup. For more information, go to Amazon Glacier.

- **Retrieving tapes** – Tapes archived to the VTS cannot be read directly. To read an archived tape, you must first retrieve it to your gateway-VTL either by using the AWS Storage Gateway console or by using the AWS Storage Gateway API. A retrieved tape will be available in your VTL in about 24 hours.

You have one VTS in your account per AWS region. For example, if you create multiple gateways in the same AWS region, one VTS will hold ejected tapes for all the gateways. If you create gateways in different AWS regions, you will have one VTS per AWS region.

After you deploy and activate a gateway-VTL, you mount the virtual tape drives and media changer on your on-premises application servers as iSCSI devices. You create virtual tapes as needed and then use your existing backup software application to write data to the virtual tapes. The media changer loads and unloads the virtual tapes into the virtual tape drives for read and write operations.

## Allocating Local Disks for the Gateway VM

Your gateway VM will need local disks, which you allocate for the following purposes:

- **Cache storage** – The cache storage acts as the durable store for data that is waiting to upload to Amazon S3 from the upload buffer.

  If your application reads data from a virtual tape, the gateway saves the data to the cache storage. The gateway stores recently accessed data in the cache storage for low-latency access. If your application requests tape data, the gateway first checks the cache storage for the data before downloading the data from AWS.
- **Upload buffer** – The upload buffer provides a staging area for the gateway before it uploads the data to a virtual tape. The upload buffer is also critical for creating recovery points that you can use to recover tapes from unexpected failures. For more information, see Recovering a Virtual Tape from a Malfunctioning Gateway-VTL (p. 282).

As your backup application writes data to your gateway, the gateway copies data to both the cache storage and the upload buffer before acknowledging completion of the write operation to your backup application.

For guidelines to determine the amount of disk space you should allocate for the cache storage and upload buffer, see Step 2.3.1: Decide the Amount of Local Disk Storage (p. 23).

# Available AWS Regions

AWS Storage Gateway stores volume, snapshot, and tape data in the AWS region in which the gateway is activated. You select a region at the upper right of the AWS Storage Gateway console before you start deploying your gateway. Following are the available regions.

| Region Name | Region String | Gateway-Cached | Gateway-Stored | Gateway-VTL |
|---|---|---|---|---|
| US East (N. Virginia) | `us-east-1` | yes | yes | yes |
| US West (N. California) | `us-west-1` | yes | yes | yes |
| US West (Oregon) | `us-west-2` | yes | yes | yes |

| Region Name | Region String | Gateway-Cached | Gateway-Stored | Gateway-VTL |
|---|---|---|---|---|
| EU (Ireland) | `eu-west-1` | yes | yes | yes |
| EU (Frankfurt) | `eu-central-1` | yes | yes | yes |
| Asia Pacific (Singapore) | `ap-southeast-1` | yes | yes | no |
| Asia Pacific (Sydney) | `ap-southeast-2` | yes | yes | yes |
| Asia Pacific (Tokyo) | `ap-northeast-1` | yes | yes | yes |
| South America (Sao Paulo) | `sa-east-1` | yes | yes | no |

# AWS Storage Gateway Pricing

For current information about pricing, go to the AWS Storage Gateway detail page.

# Accessing AWS Storage Gateway

You can use the AWS Storage Gateway console to perform various gateway configuration and management tasks. The Getting Started section and various other sections of this guide use the console to illustrate gateway functionality.

Additionally, you can use the AWS Storage Gateway API to programmatically configure and manage your gateways. For more information about the API, see API Reference for AWS Storage Gateway (p. 379).

You can also use the AWS SDKs to develop applications that interact with AWS Storage Gateway. The AWS SDKs for Java, .NET, and PHP wrap the underlying AWS Storage Gateway API to simplify your programming tasks. For information about downloading the SDK libraries, go to Sample Code Libraries.

**Next Step**

Getting Started with AWS Storage Gateway (p. 9)

# Getting Started with AWS Storage Gateway

This Getting Started section provides instructions for you to download, deploy, and activate a gateway. The steps in these instructions are common to all the storage solutions offered by AWS Storage Gateway. After completing these steps, you can proceed to one of the gateway-specific sections for additional setup instructions.

As you follow the steps in this Getting Started section, you will use the **Setup and Activate Gateway** wizard in the AWS Storage Gateway console.

Setup and Activate Gateway
close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
Choose How You Want to Run Your Gateway
Please select a configuration for running your gateway
Gateway-Cached volumes: Maintain local, low-latency access to your most recently accessed data while storing all your data in Amazon S3.
Gateway-Virtual Tape Library: Durably store data on Amazon Glacier while leveraging your existing tape-based processes.
Gateway-Stored Volumes: Schedule off-site backups to Amazon S3 for your on-premises data.

At several steps in the wizard, you perform tasks outside of the console and then return. If your session times out or the browser closes, you can return to the console.

## Recommended Reading

Before you get started, we recommend you read the following sections in this guide:

- What Is AWS Storage Gateway? (p. 1)
- How AWS Storage Gateway Works (Architecture) (p. 2)

**Next Step**

Plan Your AWS Storage Gateway Deployment (p. 10)

# Plan Your AWS Storage Gateway Deployment

The AWS Storage Gateway software appliance lets you connect your existing on-premises application infrastructure with scalable, cost-effective AWS cloud storage that provides data security features.

To deploy an AWS Storage Gateway solution, you first need to decide on the following two things:

1. **Storage solution** – Depending on your need, you can choose from one of the following storage solutions:

   - **Volume gateway (gateway-cached or gateway-stored)** – Volume gateways let you create storage volumes in the AWS cloud that your on-premises applications can access as Internet Small Computer System Interface (iSCSI) targets. There are two options—gateway-cached or gateway-stored.

     With gateway-cached volumes, you store volume data in AWS, with a small portion of recently accessed data in the cache on-premises. This approach enables low-latency access to your frequently accessed dataset and also provides seamless access to your entire dataset stored in AWS. This type of data access lets you scale your storage resource without having to provision additional hardware.

     With gateway-stored volumes, you store the entire set of volume data on-premises and store periodic point-in-time backups (snapshots) in AWS. In this model, your on-premises storage is primary, delivering low-latency access to your entire dataset, and AWS storage is the backup that you can restore in the event of a disaster in your data center.

     For an architectural overview of volume gateways, see Volume Gateway Architecture (Gateway-Cached and Gateway-Stored) (p. 2).
   - **Gateway-VTL** – If you are looking for a cost-effective, durable, long-term, off-site alternative for data archiving, you can deploy the gateway-VTL solution. The virtual tape library (VTL) interface it provides lets you leverage your existing tape-based backup software infrastructure to store data on virtual tape cartridges that you create on the gateway—for more information, see Supported Backup Software (Gateway-VTL Only) (p. 12). When you archive tapes in a virtual tape shelf (VTS), you don't worry about managing tapes on your premises and arranging shipments of tapes off-site. For an architectural overview, see Gateway–Virtual Tape Library (Gateway-VTL) Architecture (p. 5).
2. **Hosting option** – You can choose to run AWS Storage Gateway either on-premises, as a virtual machine (VM) appliance, or in AWS, as an Amazon EC2 instance. For more information, see Requirements (p. 11). If your data center goes offline and you don't have an available host, you can deploy a gateway on an EC2 instance. AWS Storage Gateway provides an Amazon Machine Image (AMI) that contains the gateway VM image.

Additionally, as you configure a host to deploy a gateway software appliance, you will need to allocate sufficient storage for the gateway VM.

Before you continue to the next step, make sure you have done the following:

1. For a gateway deployed on-premises, you have decided the type of host you want to set up (VMware ESXi Hypervisor or Microsoft Hyper-V 2008 R2) and set it up. If you deploy the gateway behind a firewall, you must make sure certain ports are accessible to the gateway VM. For more information, see Requirements (p. 11). The following topics provide steps for configuring the host:
   - Appendix B: Configuring a VMware ESXi Host for AWS Storage Gateway (p. 406)
   - Appendix D: Configuring a Microsoft Hyper-V Host for AWS Storage Gateway (p. 411)
2. For a gateway-VTL, you have installed client backup software. For more information, Supported Backup Software (Gateway-VTL Only) (p. 12).

**Next Step**

Requirements (p. 11)

# Requirements

Unless otherwise noted, the following requirements are common to all gateway configurations.

**Topics**


- Supported Hypervisors and Host Requirements (p. 11)
- Hardware Requirements (p. 11)
- Supported iSCSI Initiators (p. 12)
- Supported Backup Software (Gateway-VTL Only) (p. 12)
- Storage Requirements (p. 12)
- Network Requirements (p. 12)


## Supported Hypervisors and Host Requirements

You can choose to run AWS Storage Gateway either on-premises as a virtual machine (VM) appliance, or in AWS as an Amazon Elastic Compute Cloud (Amazon EC2) instance.

AWS Storage Gateway supports the following hypervisor versions and hosts:

- VMware ESXi Hypervisor (version 4.1, 5.0, 5.1, or 5.5)—A free version of VMware is available on the VMware website. You will also need a VMware vSphere client to connect to the host.
- Microsoft Hyper-V Hypervisor (version 2008 R2, 2012, or 2012 R2)—A free, stand-alone version of Hyper-V is available at the Microsoft Download Center. You will need a Microsoft Hyper-V Manager on a Microsoft Windows client computer to connect to the host.
- EC2 instance—AWS Storage Gateway provides an Amazon Machine Image (AMI) that contains the gateway VM image. Only gateway-cached volumes and gateway-VTLs can be deployed on Amazon EC2. For information about how to deploy a gateway on Amazon EC2, see Amazon EC2 Gateway (p. 57).

## Hardware Requirements

When deploying your gateway on-premises, you must make sure that the underlying hardware on which you are deploying the gateway VM is able to dedicate the following resources:

- Four or eight virtual processors assigned to the VM.

  For more information, see Optimizing Gateway Performance (p. 158) or Optimizing Gateway-VTL Performance (p. 341).
- 7.5 GB of RAM assigned to the VM
- 75 GB of disk space for installation of VM image and system data

When deploying your gateway on Amazon EC2, you must use the **m** and **c1** instance types and the instance size must be at least size **xlarge**. You must select one of these instance types for the gateway to function. For more information, go to AWS Storage Gateway in AWS Marketplace.

# Supported iSCSI Initiators

When you deploy a gateway-cached or gateway-stored volume gateway, you can create iSCSI storage volumes on your gateway. When you deploy a gateway-VTL, the gateway is preconfigured with one media changer and ten tape drives. These tape drives and the media changer are available to your existing client backup applications as iSCSI devices. To connect to these iSCSI devices, AWS Storage Gateway supports the following iSCSI initiators:

- Windows Server 2012 and Windows Server 2012 R2
- Windows Server 2008 and Windows Server 2008 R2
- Windows 7
- Red Hat Enterprise Linux 5
- Red Hat Enterprise Linux 6
- VMware ESX Initiator

The VMware ESX Initiator provides an alternative to using initiators in the guest operating systems of your VMs.

# Supported Backup Software (Gateway-VTL Only)

Typically, you will use a backup application to read, write, and manage tapes with a gateway-VTL.

The following lists the third-party backup software that Gateway-VTL supports.

- Symantec NetBackup version 7.x
- Symantec Backup Exec 2012
- Symantec Backup Exec 2014
- Symantec Backup Exec 15
- Microsoft System Center 2012 R2 Data Protection Manager
- Veeam Backup & Replication V7
- Veeam Backup & Replication V8
- Dell NetVault Backup 10.0

# Storage Requirements

In addition to 75 GB disk space for the VM, you will also need additional disks for the gateway.

- For gateway-cached volume configuration, you will need storage for the local cache and an upload buffer.
- For gateway-stored volume configuration, you will need storage for your entire dataset and an upload buffer.
- For gateway-VTL configuration, you will need storage for the local cache and an upload buffer.

For more information about how to add disks, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23).

# Network Requirements

AWS Storage Gateway requires the following TCP ports for its operation.

| Protocol | Port | Source | Destination | How Used |
|---|---|---|---|---|
| TCP | 443 | Storage Gateway | Internet | For communication from AWS Storage Gateway to the AWS service endpoint. |
| TCP | 80 | Local Networks | Storage Gateway | By local systems to obtain the storage gateway activation key. Port 80 is only used during activation of the Storage Gateway appliance.<br>**Note**<br>AWS Storage Gateway does not require port 80 to be publicly accessible. The required level of access to port 80 depends on your network configuration. If you activate your gateway from the AWS Storage Gateway Management Console, the host from which you connect to the console must have access to your gateway's port 80. |

| Protocol | Port | Source | Destination | How Used |
|---|---|---|---|---|
| TCP | 3260 | Local Networks | Storage Gateway | By local systems to connect to iSCSI targets exposed by the gateway. |

**Next Step**

Step 1: Sign Up for AWS Storage Gateway (p. 14)

# Step 1: Sign Up for AWS Storage Gateway

To use AWS Storage Gateway, you need an AWS account that gives you access to all AWS resources, forums, support, and usage reports. You are not charged for any of the services unless you use them. If you already have an AWS account, you can skip this step.

### To sign up for AWS account

1. Open http://www.amazonaws.cn/, and then click **Sign Up**.
2. Follow the on-screen instructions.

   Part of the sign-up procedure involves receiving a phone call and entering a PIN using the phone keypad.

**Next Step**

Step 2: Deploy a VM and Activate the Gateway (p. 14)

# Step 2: Deploy a VM and Activate the Gateway

Now that you have decided on the hosting option, you deploy and activate your gateway.

If you are deploying the gateway on-premises, you download and deploy the gateway virtual machine (VM) and then activate the gateway. If you are deploying the gateway on an Amazon EC2 instance, you launch an Amazon Machine Image (AMI) that contains the gateway VM image and then activate the gateway.

### To deploy a gateway on premises

1. Go to the AWS Console at http://console.www.amazonaws.cn/storagegateway.

2. In the Navigation pane, click **Deploy a new Gateway**.

### To deploy a gateway on Amazon EC2

1. Go to the AWS Console at http://console.www.amazonaws.cn/storagegateway.
2. On the **Navigation** pane, click **Deploy a new Gateway on Amazon EC2**.

### Next Step

Depending on whether you are deploying your gateway on premises or on Amazon EC2, click one of the following links.

- On-Premises Gateway (p. 15)
- Amazon EC2 Gateway (p. 57)

# On-Premises Gateway

When you select **Deploy a New Gateway** to deploy a gateway on-premises, the setup wizard asks you to choose a gateway configuration.

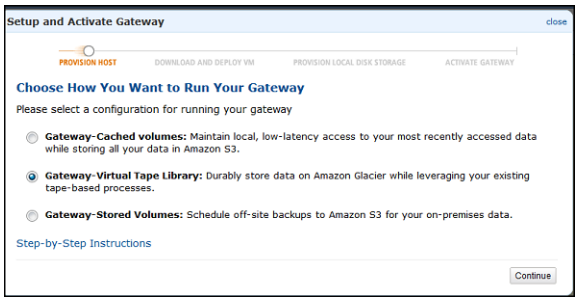
Setup and Activate Gateway
close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
Choose How You Want to Run Your Gateway
Please select a configuration for running your gateway
Gateway-Cached volumes: Maintain local, low-latency access to your most recently accessed data while storing all your data in Amazon S3.
Gateway-Virtual Tape Library: Durably store data on Amazon Glacier while leveraging your existing tape-based processes.
Gateway-Stored Volumes: Schedule off-site backups to Amazon S3 for your on-premises data.
Step-by-Step Instructions
Continue

### To choose a gateway configuration

1. Select the gateway you want to set up.
2. Click **Continue**.

### Next Step

Depending on the hypervisor host you provisioned in your data center, click one the following links.

- Step 2.1: Provision a VMware Host to Deploy the AWS Storage Gateway VM (p. 16)
- Step 2.1: Provision a Hyper-V Host to Deploy the AWS Storage Gateway VM (p. 35)

# VMware Host

This section shows you how to perform the tasks described following.

### Topics

- Step 2.1: Provision a VMware Host to Deploy the AWS Storage Gateway VM (p. 16)
- Step 2.2: Download and Deploy the AWS Storage Gateway VM on Your Host (p. 16)
- Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23)
- Step 2.4: Activate the AWS Storage Gateway (p. 30)

- Using AWS Storage Gateway with VMware High Availability (p. 35)

**Next Step**

Step 2.1: Provision a VMware Host to Deploy the AWS Storage Gateway VM (p. 16)

## Step 2.1: Provision a VMware Host to Deploy the AWS Storage Gateway VM

If you have not provisioned a host to deploy the gateway VM, provision a host now.

Setup and Activate Gateway
close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
You will need a host in your datacenter to deploy the gateway virtual machine (VM). Pick a host that meets these minimum requirements.
« Back
Continue

### To provision a host

1. Review the minimum host requirements in Requirements (p. 11).
2. Set up a host in your data center with the VMware ESXi hypervisor. For a minimum set of instructions to install the hypervisor, see Appendix B: Configuring a VMware ESXi Host for AWS Storage Gateway (p. 406).

   **Note**
   If you plan to deploy AWS Storage Gateway using VMware High Availability (HA) for failover protection, see Using AWS Storage Gateway with VMware High Availability (p. 35). In this tutorial exercise, you deploy your AWS Storage Gateway VM on a single host with no clustering or failover provision.
3. In the setup wizard, click **Continue**.

**Next Step**

Step 2.2: Download and Deploy the AWS Storage Gateway VM on Your Host (p. 16)

## Step 2.2: Download and Deploy the AWS Storage Gateway VM on Your Host

The AWS Storage Gateway virtual machine is available as a VMware ESX `.ova` package. This section shows you how to perform the tasks described following.

**Topics**

- Step 2.2.1: Download the AWS Storage Gateway VM (p. 16)
- Step 2.2.2: Deploy the AWS Storage Gateway VM to Your Host (p. 17)
- Step 2.2.3: Synchronize VM Time with Host Time (p. 20)

### Step 2.2.1: Download the AWS Storage Gateway VM

### To download the VM

1. In the AWS Storage Gateway console's **Setup and Activate Gateway** wizard, navigate to the **DOWNLOAD AND DEPLOY VM** page.
2. Select **I want to run the AWS Storage Gateway on VMware ESXi**, and then click **Continue**.

Setup and Activate Gateway
Close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
Select a virtualization platform to run the AWS Storage Gateway Virtual Machine (VM).
I want to run the AWS Storage Gateway on VMware ESXi
I want to run the AWS Storage Gateway on Microsoft Hyper-V 2008 R2
I want to run the AWS Storage Gateway on Microsoft Hyper-V 2012
« Back
Continue

3. Click **Download** to download a `.zip` file that contains the `.ova` file. Extract and save the `.zip` file to a location on your computer.

   **Note**
   The `.zip` file is over 500 MB in size and may take some time to download, depending on your network connection.

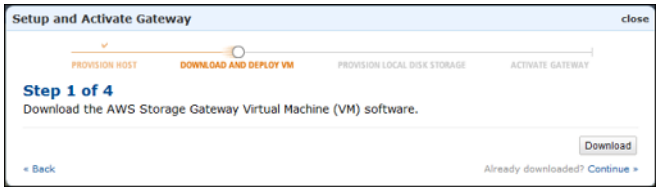


## Step 2.2.2: Deploy the AWS Storage Gateway VM to Your Host

### To deploy the AWS Storage Gateway VM to your host

**Note**
Depending on the version of hypervisor you use, the graphical user interface (GUI) shown in these procedures may look slightly different. For more information, see Supported Hypervisors and Host Requirements (p. 11).

1. To connect to your hypervisor host, start the VMware vSphere client on Windows, type your host's IP and your credentials in the login dialog box, and then click **Login**.
2. Deploy the AWS Storage Gateway VM on the host using the following steps.

   a. On the **File** menu of the vSphere client, click **Deploy OVF Template**.

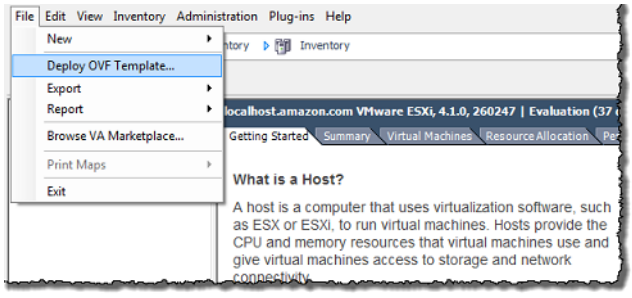


   Doing this opens the **Deploy OVF Template** wizard, a series of steps to provide information required to deploy the VM.

   b. In the **Source** pane, provide the file path to the AWS Storage Gateway `.ova` package, and then click **Next**.

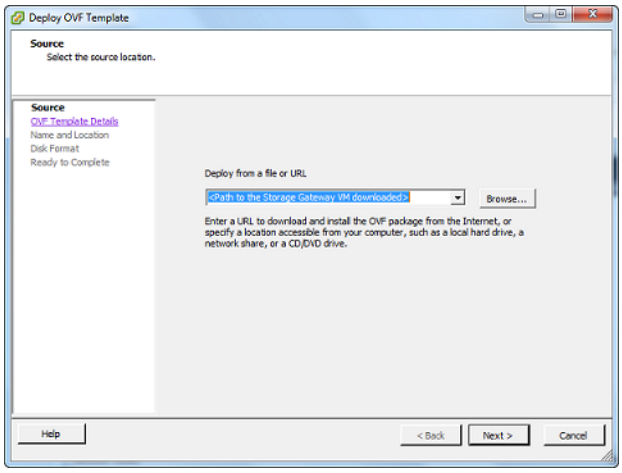


c. In the **OVF Template Details** pane, click **Next**.

d. In the **Name and Location** pane, enter the VM name in the **Name** field, and then click **Next**.

   Although the vSphere client displays this VM name, AWS Storage Gateway does not use this name.

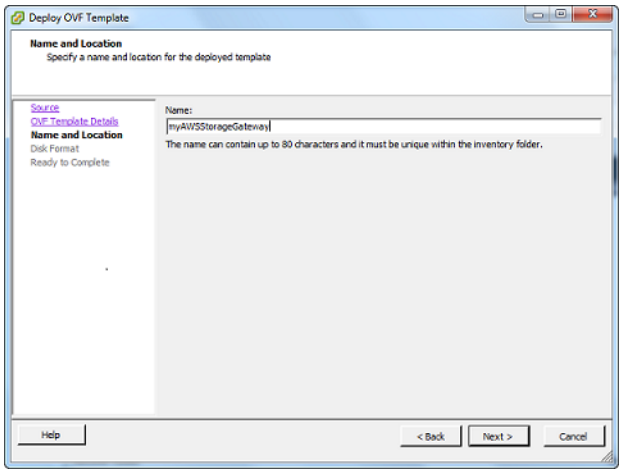


e. The following **Datastore** pane appears only if your host has multiple data stores. In this pane, you select the data store where you want to deploy the VM, and then click **Next**. If your host has only one data store, skip to the next step.

   A *data store* is a virtual representation of underlying physical storage resources. The following example shows a host that has two data stores: datastore1 and datastore2.

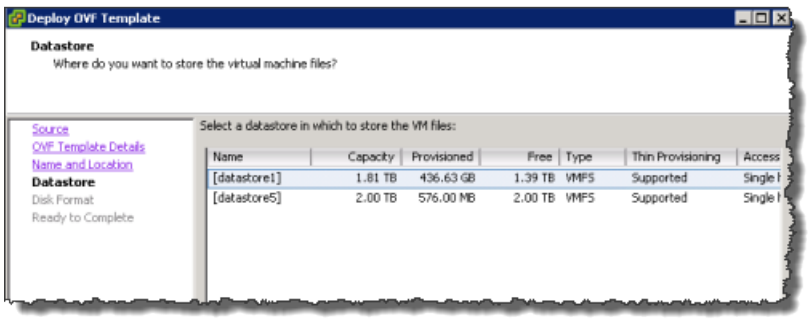


f. In the **Disk Format** pane, select **Thick provisioned format**, and then click **Next**.

When you use thick provisioning, the disk storage is allocated immediately, resulting in better performance. In contrast, thin provisioning allocates storage on demand. On-demand allocation can affect the normal functioning of AWS Storage Gateway. For AWS Storage Gateway to function properly, the VM disks must be stored in thick-provisioned format.

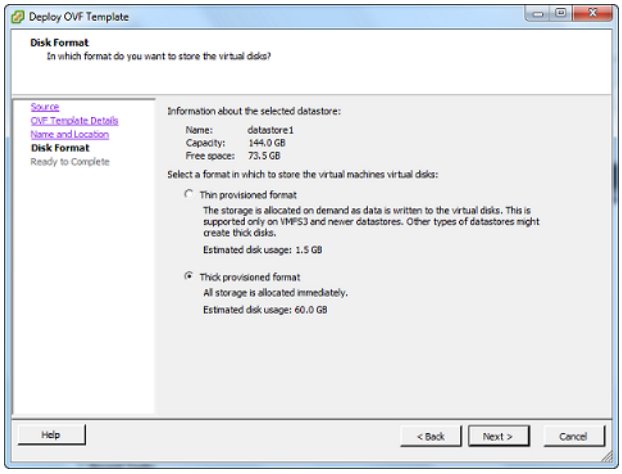
Deploy OVF Template
Disk Format
In which format do you want to store the virtual disks?
Source
OVF Template Details
Name and Location
Disk Format
Ready to Complete
Information about the selected datastore:
Name: datastore1
Capacity: 144.0 GB
Free space: 73.5 GB
Select a format in which to store the virtual machines virtual disks:
Thin provisioned format
The storage is allocated on demand as data is written to the virtual disks. This is supported only on VMFS3 and newer datastores. Other types of datastores might create thick disks.
Estimated disk usage: 1.5 GB
Thick provisioned format
All storage is allocated immediately.
Estimated disk usage: 60.0 GB
Help
< Back
Next >
Cancel

g. In the **Ready to Complete** pane, click **Finish**.

The AWS Storage Gateway VM starts deploying to your host.

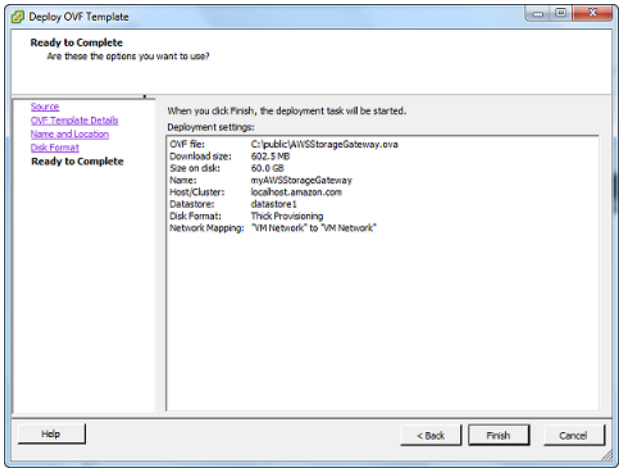


h. View the details of the new VM to verify the information about the VM.

   i. Depending on the state of your vSphere client, you might need to click the **Inventory** icon to view the host object that contains the new VM.

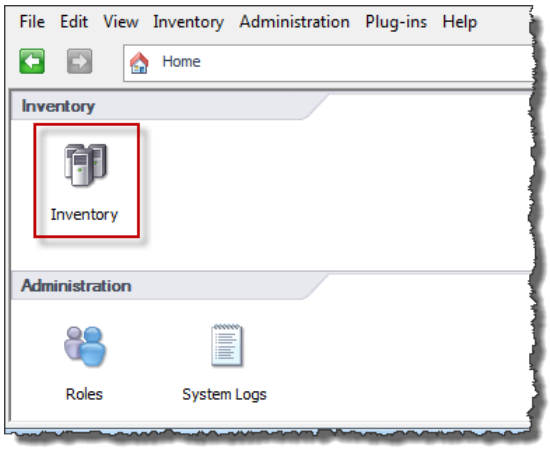


ii. Expand the host object to view the details of the new VM.

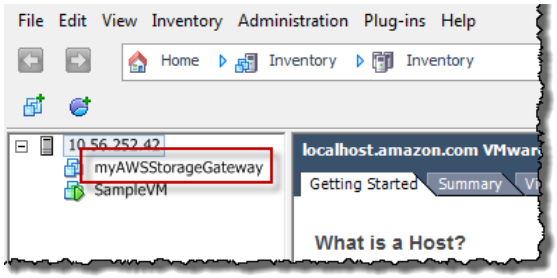


## Step 2.2.3: Synchronize VM Time with Host Time

To successfully activate your gateway, you must ensure that your VM time is synchronized to the host time, and that the host time is correctly set. In this section, you first synchronize the time on the VM to the host time. Then you check the host time and, if needed, set the host time and configure the host to synchronize its time automatically to a Network Time Protocol (NTP) server.

**Important**
Synchronizing the VM time with the host time is required for successful gateway activation.

### To synchronize VM time with host time

1. Configure your VM time.

   a. In the vSphere client, right-click the name of your gateway VM and click **Edit Settings**.

      The **Virtual Machine Properties** dialog box opens.

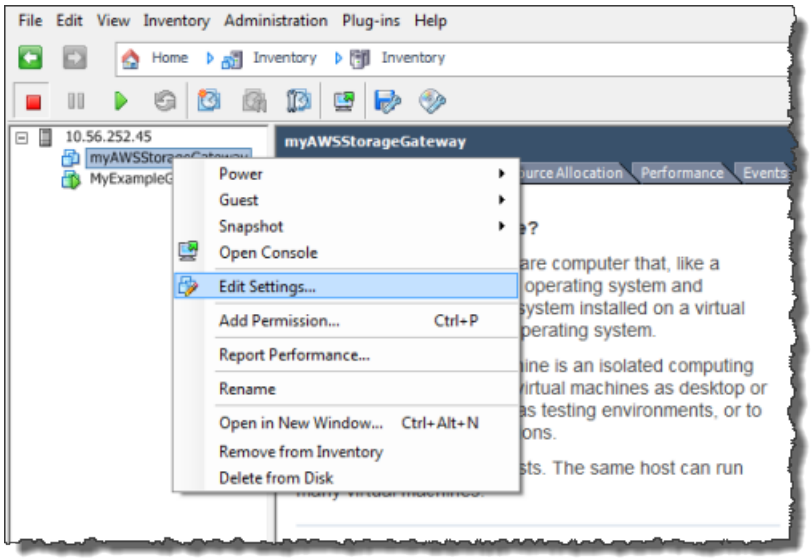


b. Click the **Options** tab, and click **VMware Tools** in the options list.

c. Check the **Synchronize guest time with host** option, and then click **OK**.

The VM synchronizes its time with the host.

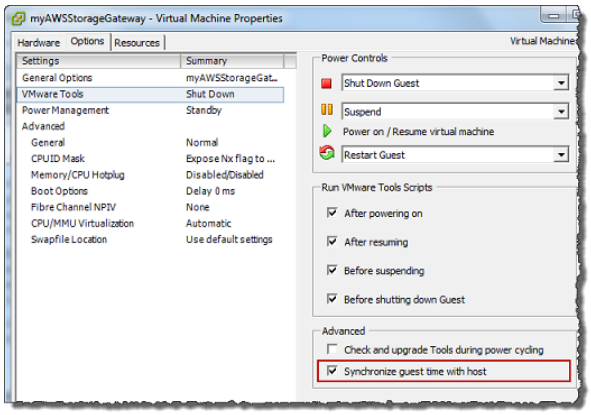


2. Configure the host time.

It is important to make sure that your host clock is set to the correct time. If you have not configured your host clock, perform the following steps to set and synchronize it with an NTP server.

a. In the VMware vSphere client, select the vSphere host node in the left pane, and then click the **Configuration** tab.

b. Select **Time Configuration** in the **Software** panel, and then click the **Properties** link.

The **Time Configuration** dialog box appears.

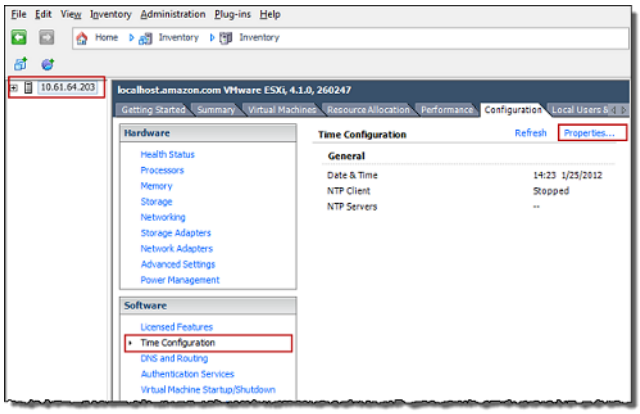


c. In the **Date and Time** panel, set the date and time.

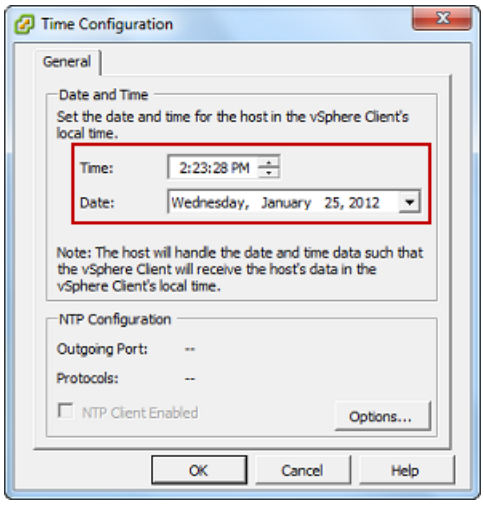
Time Configuration
General
Date and Time
Set the date and time for the host in the vSphere Client's local time.
Time: 2:23:28 PM
Date: Wednesday, January 25, 2012
Note: The host will handle the date and time data such that the vSphere Client will receive the host's data in the vSphere Client's local time.
NTP Configuration
Outgoing Port: --
Protocols: --
NTP Client Enabled
Options...
OK
Cancel
Help

d. Configure the host to synchronize its time automatically to an NTP server.

   i. Click **Options** in the **Time Configuration** dialog box, and then in the **NTP Daemon (ntpd) Options** dialog box, click **NTP Settings** in the left pane.

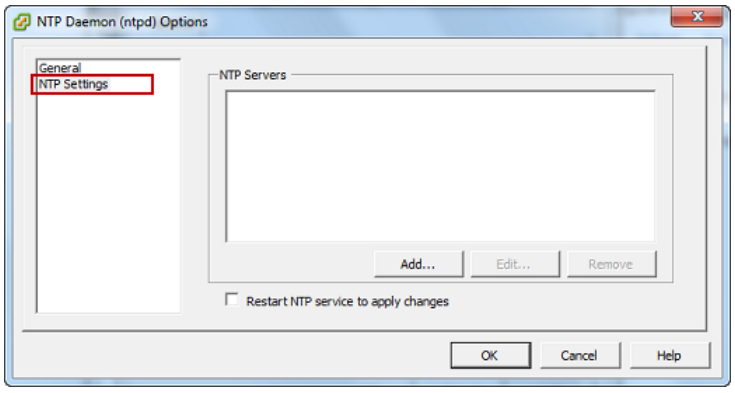


   ii. Click **Add** to add a new NTP server.

   iii. In the **Add NTP Server** dialog box, type the IP address or the fully qualified domain name of an NTP server, and then click **OK**.

   You can use `pool.ntp.org` as shown in the following example.

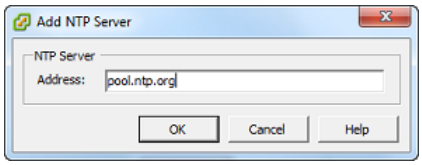


iv. In the **NTP Daemon (ntpd) Options** dialog box, click **General** in the left pane.

v. In the **Service Commands** pane, click **Start** to start the service.

Note that if you change this NTP server reference or add another later, you will need to restart the service to use the new server.

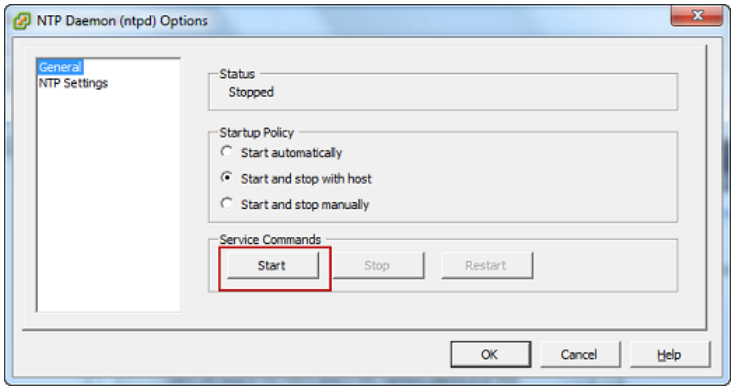


e. Click **OK** to close the **NTP Daemon (ntpd) Options** dialog box.

f. Click **OK** to close the **Time Configuration** dialog box.

**Next Step**

Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23)

## Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare)

Next, you allocate local storage for your deployed gateway VM. After you activate your gateway, you configure this local storage for the gateway's use.

**Topics**


- Step 2.3.1: Decide the Amount of Local Disk Storage (p. 23)
- Step 2.3.2: Provision Local Disk Storage (p. 24)
- Step 2.3.3: Configure the AWS Storage Gateway VM to Use Paravirtualized Disk Controllers (p. 29)


### Step 2.3.1: Decide the Amount of Local Disk Storage

In this step, you decide the number and size of disks to allocate for your gateway, as follows:

- All gateways require one or more disks for an upload buffer.
- Depending on the storage solution you deploy (see Plan Your AWS Storage Gateway Deployment (p. 10)), the gateway requires the following additional storage:
  - Gateway-VTL and gateway-cached volumes require one or more disks for cache storage.

- Gateway-stored volumes require one or more disks for volume storage—one local disk required for each volume you will create on the gateway.

You need at least two disks to begin with. The following table recommends sizes for local disk storage for your deployed gateway.

| Disk Size | Gateway-Cached | Gateway-Stored | Gateway-VTL |
|---|---|---|---|
| Upload buffer (minimum) | 150 GiB | 150 GiB | 150 GiB |
| Upload buffer (maximum) | 2 TiB | 2 TiB | 2 TiB |
| Cache storage (minimum) | 1.1 x upload buffer | — | 1.1 x upload buffer |
| Cache storage (maximum) | 16 TiB | — | 16 TiB |
| Volume storage (minimum) | — | 1 GiB | — |
| Volume storage (maximum) | 32 TiB | 1 TiB | — |

Before going to the next step, decide the number and size of disks to allocate. You can add more local storage after you set the gateway up, and as your workload demands. If you're not sure what number and size of disks to use, provision two local disks with the minimum recommended size for your gateway type.

If you plan to deploy your gateway in production, you should consider your real workload in determining disk sizes. For disk size guidelines, see one of the following sections depending on the type of gateway you are deploying:

- Gateway-cached: Managing the Upload Buffer and Cache Storage (Gateway-Cached) (p. 169).
- Gateway-stored: Managing the Upload Buffer (Gateway-Stored) (p. 178)
- Gateway-VTL: Managing the Upload Buffer and Cache Storage (Gateway-VTL) (p. 304)

For more information about how gateways use local storage, see How AWS Storage Gateway Works (Architecture) (p. 2).

In the next step, you allocate the local disk storage to the gateway VM you deployed.

### Step 2.3.2: Provision Local Disk Storage

You can add disks from direct attached storage (DAS) or from a storage area network (SAN). This section tells how to add a virtual disk from a DAS. For instructions on attaching iSCSI volumes from an existing SAN for this step, see To add a new iSCSI target (p. 407).

Repeat the following procedure for each local disk you planed in the preceding step.

#### To allocate a local disk from DAS

1. Start the VMware vSphere client, and then connect to your host.
2. In the client, right-click the name of your gateway VM, and then click **Edit Settings**.

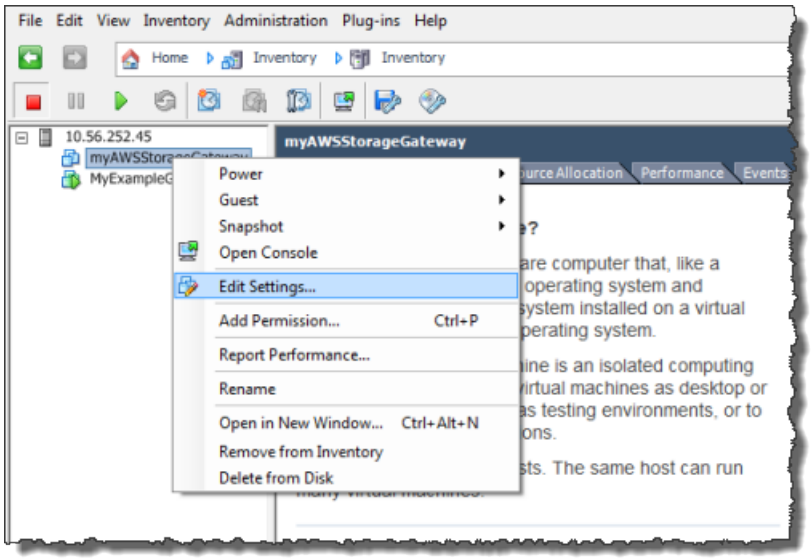


3. Click the **Hardware** tab of the **Virtual Machine Properties** dialog box, and then click **Add** to add a device.

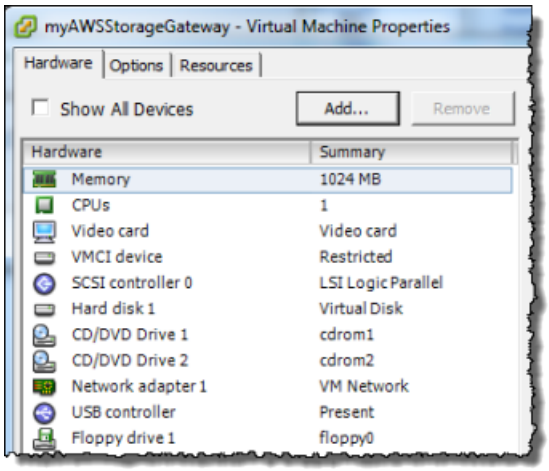


4. Follow the steps in the **Add Hardware** wizard to add a disk:

   a. In the **Device Type** pane, click **Hard Disk** to add a disk, and then click **Next**.

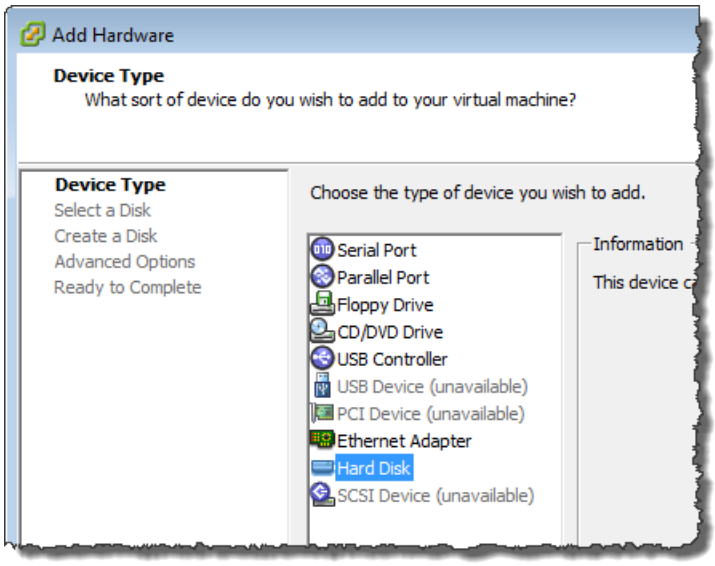


b. In the **Select a Disk** pane, select **Create a new virtual disk**, and then click **Next**.

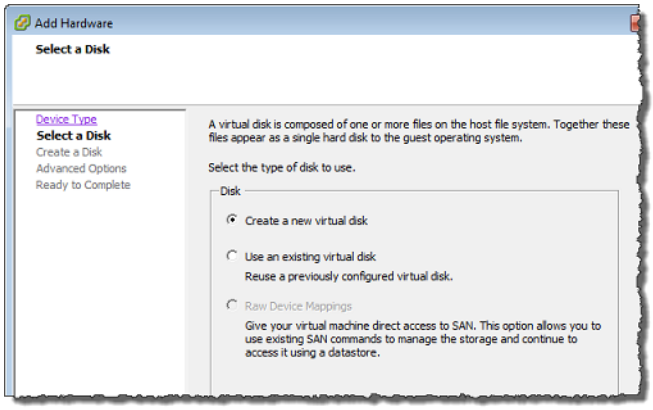
Add Hardware
Select a Disk
Device Type
Select a Disk
Create a Disk
Advanced Options
Ready to Complete
A virtual disk is composed of one or more files on the host file system. Together these files appear as a single hard disk to the guest operating system.
Select the type of disk to use.
Disk
Create a new virtual disk
Use an existing virtual disk
Reuse a previously configured virtual disk.
Raw Device Mappings
Give your virtual machine direct access to SAN. This option allows you to use existing SAN commands to manage the storage and continue to access it using a datastore.

c. In the **Create a Disk** pane, specify the following.

- Capacity: The disk size you previously decided.

  **Note**
  The disk sizes used in these examples are not suitable for real-world workloads. If your setup is a trial setup, you can use small disk sizes.

- Disk provisioning: We recommend thick provisioning,
- Location: We strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store different from the one used for the gateway VM. In addition, we strongly recommend you dedicate a data store for the disks that you will use for upload buffer and another data store for other disks (see Step 2.3.1: Decide the Amount of Local Disk Storage (p. 23)). For this example, we store with the VM.

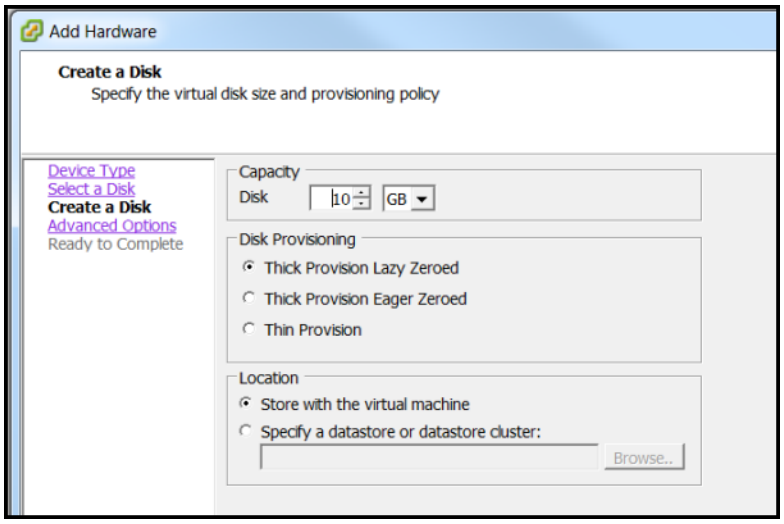


d. Click **Next**, and in the **Advanced Options** pane, accept the default values, and then click **Next**.

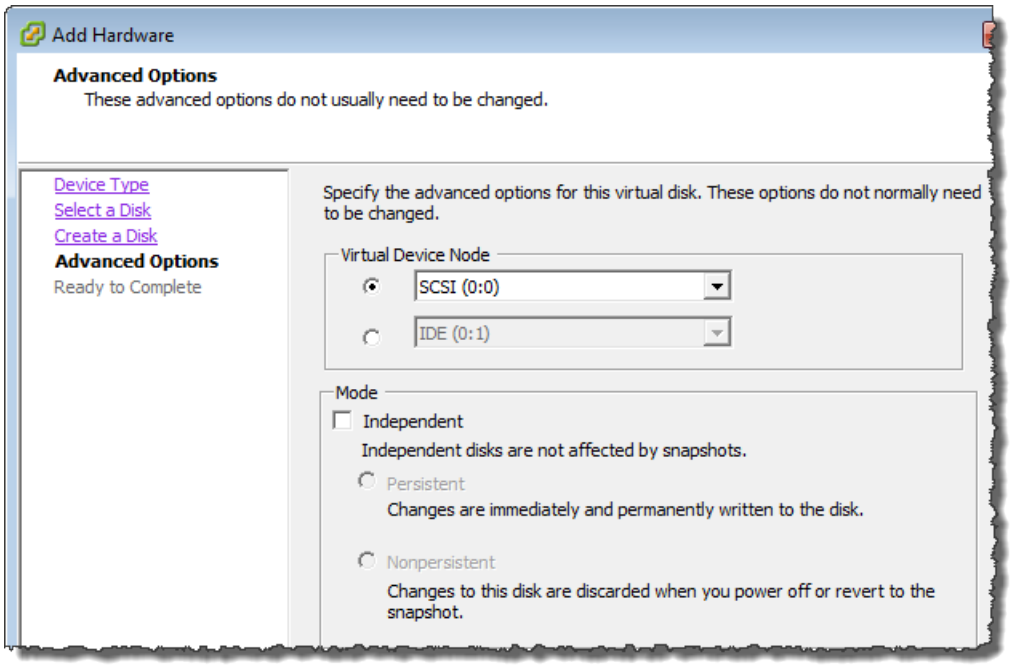
Add Hardware
Advanced Options
These advanced options do not usually need to be changed.
Device Type
Select a Disk
Create a Disk
Advanced Options
Ready to Complete
Specify the advanced options for this virtual disk. These options do not normally need to be changed.
Virtual Device Node
SCSI (0:0)
IDE (0:1)
Mode
Independent
Independent disks are not affected by snapshots.
Persistent
Changes are immediately and permanently written to the disk.
Nonpersistent
Changes to this disk are discarded when you power off or revert to the snapshot.

e. In the **Ready to Complete** pane, accept the default values, and then click **Finish**.

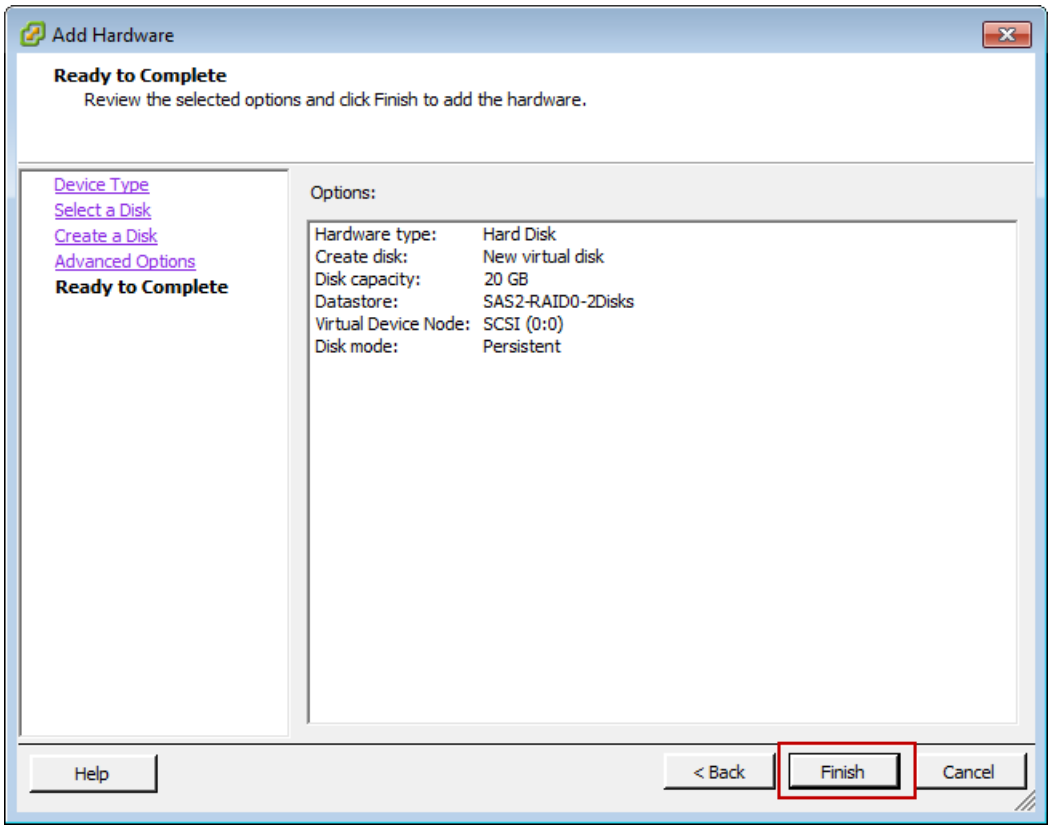


f. In the **Virtual Machine Properties** dialog box, click **OK** to complete adding the disk.

5. Repeat steps 1 through 4 for each disk you planned to allocate.
6. Verify your gateway VM has the disks you added by following these steps:

   a. In the client, right-click the name of your gateway VM and click **Edit Settings**.
   b. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, verify the disks in the hardware list.

      For example, the following screenshot shows two disks added to the gateway. The disks you add appear in the AWS Storage Gateway console as SCSI (0:0), SCSI (0:1), and so on.

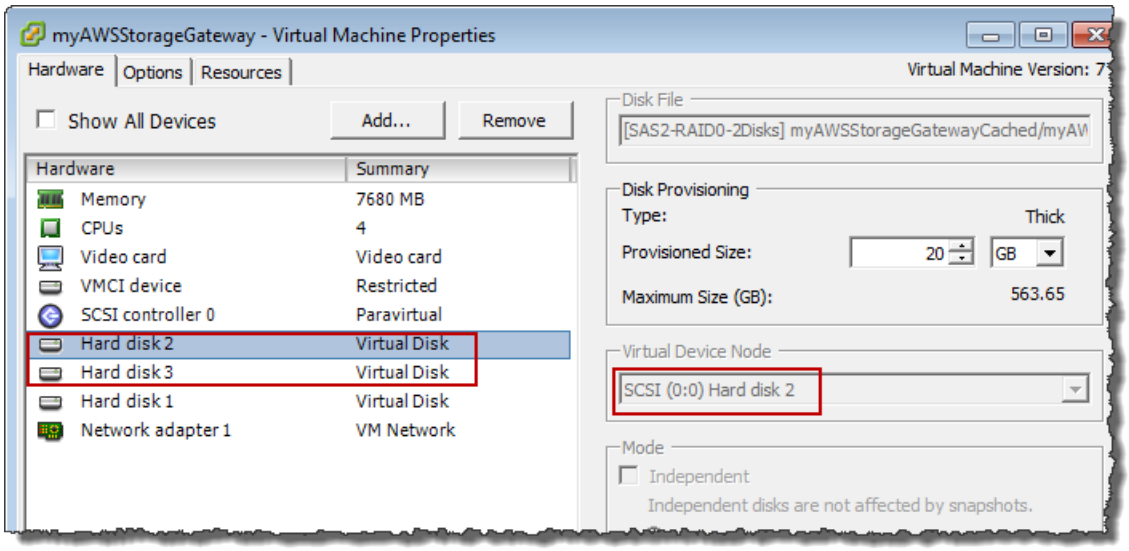


## Step 2.3.3: Configure the AWS Storage Gateway VM to Use Paravirtualized Disk Controllers

In this task, you set the iSCSI controller so that the VM uses paravirtualization. *Paravirtualization* is a mode where the gateway VM works with the host operating system so the console can identify the virtual disks that you add to your VM.

> **Note**
> You must complete this step to avoid issues in identifying these disks when you configure them in the gateway console.

### To configure your VM to use paravirtualized controllers

1. In the VMware vSphere client, right-click the name of your gateway VM, and then click **Edit Settings**.
2. In the **Virtual Machine Properties** dialog box, click the **Hardware** tab, select the **SCSI controller 0**, and then click **Change Type**.

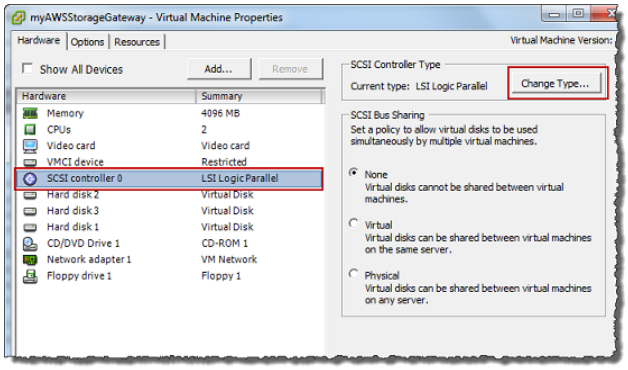


3. In the **Change SCSI Controller Type** dialog box, select the **VMware Paravirtual** SCSI controller type, and then click **OK**.

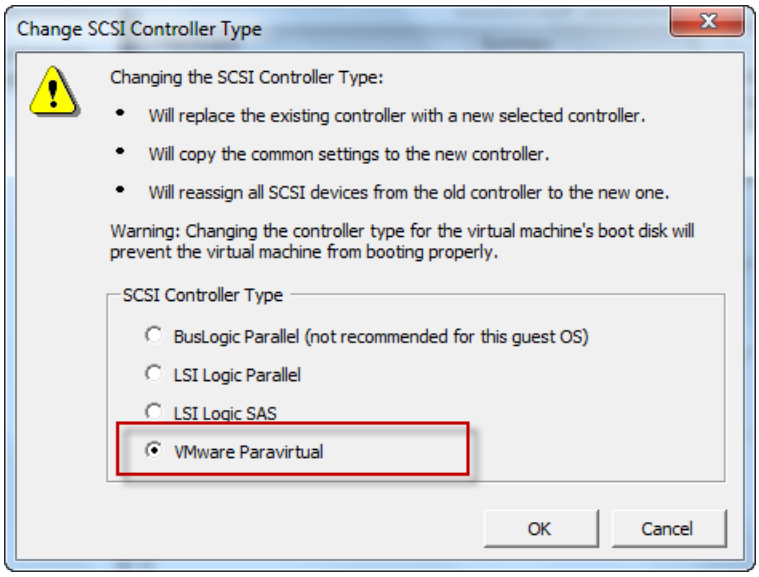
Change SCSI Controller Type
Changing the SCSI Controller Type:
Will replace the existing controller with a new selected controller.
Will copy the common settings to the new controller.
Will reassign all SCSI devices from the old controller to the new one.
Warning: Changing the controller type for the virtual machine's boot disk will prevent the virtual machine from booting properly.
SCSI Controller Type
BusLogic Parallel (not recommended for this guest OS)
LSI Logic Parallel
LSI Logic SAS
VMware Paravirtual
OK
Cancel

**Next Step**

Step 2.4: Activate the AWS Storage Gateway (p. 30)

## Step 2.4: Activate the AWS Storage Gateway

Now, you can activate your gateway. The activation process associates your gateway with your AWS account. You must turn on the gateway VM before you activate your gateway.

### To activate your gateway

1. Turn on the VM.

   a. In the vSphere client, select the gateway VM.

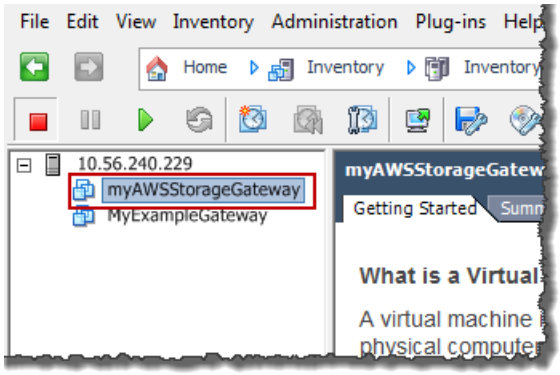


   b. On the **Toolbar** menu, click the **Power On** icon.

      Your gateway VM icon now has a green arrow associated with it, indicating that you have turned on the VM.

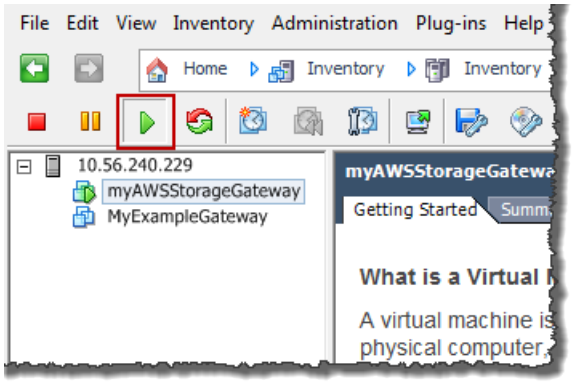


2. Activate your gateway.

   a. Obtain the IP address of your gateway.

      i. In the vSphere client, select your deployed gateway VM.

      ii. Click the **Summary** tab and take note of the IP address (you might want to write it down). You will use this IP address to activate the gateway.

      **Note**
      The IP address of your gateway appears on the **Summary** tab. After turning on the VM, it might take a few moments for the IP address to appear.

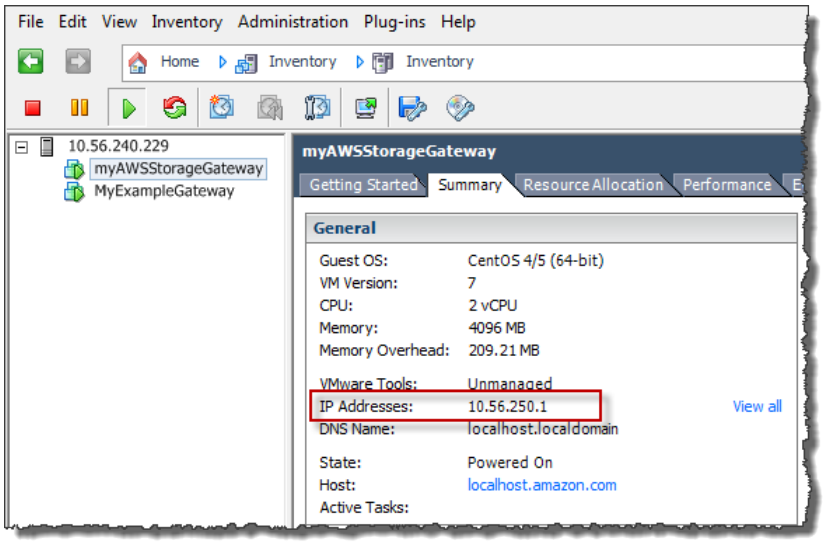


   b. Associate your gateway with your AWS account.

      i. In the AWS Storage Gateway console's **Setup and Activate Gateway** wizard, navigate to the **ACTIVATE GATEWAY** page.

      If the wizard has not already started, click the **Set up and Activate a New Gateway** button, and then click **Continue** in each wizard step until you reach that page.

      ii. In the **Enter IP address** text box, type the IP address of your gateway, and then click **Proceed to Activation**.

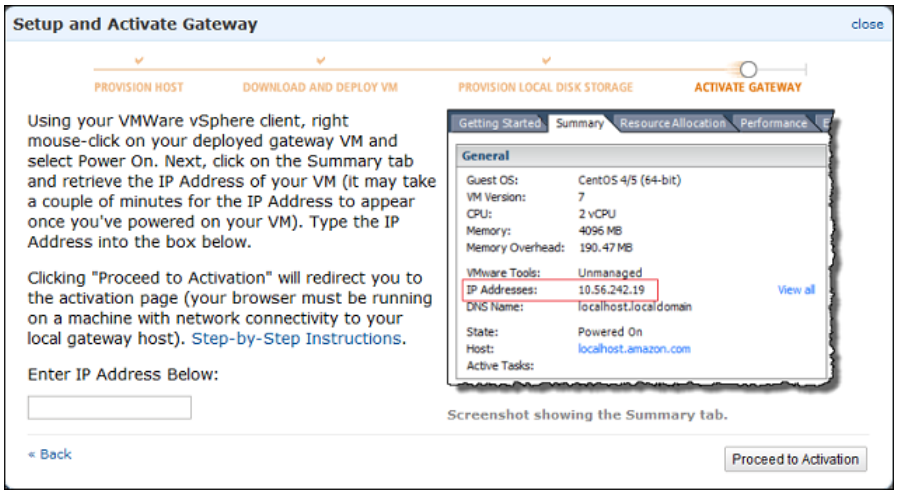
Setup and Activate Gateway
close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
Using your VMWare vSphere client, right mouse-click on your deployed gateway VM and select Power On. Next, click on the Summary tab and retrieve the IP Address of your VM (it may take a couple of minutes for the IP Address to appear once you've powered on your VM). Type the IP Address into the box below.
Clicking "Proceed to Activation" will redirect you to the activation page (your browser must be running on a machine with network connectivity to your local gateway host). Step-by-Step Instructions.
Enter IP Address Below:
Getting Started
Summary
Resource Allocation
Performance
General
Guest OS: CentOS 4/5 (64-bit)
VM Version: 7
CPU: 2 vCPU
Memory: 4096 MB
Memory Overhead: 190.47 MB
VMware Tools: Unmanaged
IP Addresses: 10.56.242.19
DNS Name: localhost.localdomain
View all
State: Powered On
Host: localhost.amazon.com
Active Tasks:
Screenshot showing the Summary tab.
« Back
Proceed to Activation

**Note**

During activation, your browser connects to the gateway. If activation fails, check that the IP address you entered is correct. If the IP address is correct, confirm that your network is configured to let your browser access the gateway VM.

iii. To complete the activation process, enter the information listed following on the activation page. The activation page you see depends on your gateway configuration.

- **AWS Region** determines where AWS stores your snapshots (for gateway-cached volumes) or virtual tapes (for gateway-VTLs). This region is the one in which you deployed your gateway.
- **Gateway Time Zone** specifies the time zone to use for your gateway.
- **Gateway Name** identifies your gateway. You use this name to manage your gateway in the console; you can change it after the gateway is activated. This name must be unique to your account.
- **Medium Changer Type** specifies the type of medium changer to use for your backup software (available for gateway-VTLs only).

**Important**

The type of medium changer you select depends on the backup software you plan to use. The following table shows which medium changer to select for your backup software.

| Backup Software | Medium Changer Type |
|---|---|
| Symantec NetBackup Version 7.x | `AWS-Gateway-VTL` |
| Symantec Backup Exec 2012 | `STK-L700` |
| Symantec Backup Exec 2014 | `AWS-Gateway-VTL` |
| Symantec Backup Exec 15 | `AWS-Gateway-VTL` |
| Microsoft System Center 2012 R2 Data Protection Manager | `STK-L700` |
| Veeam Backup & Replication V7 | `STK-L700` |
| Veeam Backup & Replication V8 | `STK-L700` |
| Dell NetVault Backup 10.0 | `STK-L700` |

**Note**
You must select the medium changer that is recommended for your backup software. Other medium changers might not function properly. You can select a different medium changer after the gateway is activated. For more information, see Selecting a Medium Changer After Gateway Activation (p. 424).

- **Tape Drive Type** specifies the type of tape drive used by this gateway (available for gateway-VTLs only).

**Note**
AWS Storage Gateway is supported in the following regions.

| Region Name | Region String | Gateway-Cached | Gateway-Stored | Gateway-VTL |
|---|---|---|---|---|
| US East (N. Virginia) | `us-east-1` | yes | yes | yes |
| US West (N. California) | `us-west-1` | yes | yes | yes |
| US West (Oregon) | `us-west-2` | yes | yes | yes |
| EU (Ireland) | `eu-west-1` | yes | yes | yes |
| EU (Frankfurt) | `eu-central-1` | yes | yes | yes |
| Asia Pacific (Singapore) | `ap-south-east-1` | yes | yes | no |
| Asia Pacific (Sydney) | `ap-south-east-2` | yes | yes | yes |
| Asia Pacific (Tokyo) | `ap-north-east-1` | yes | yes | yes |
| South America (Sao Paulo) | `sa-east-1` | yes | yes | no |

The following screenshot shows the activation page for gateway-cached volumes or gateway-stored volumes.

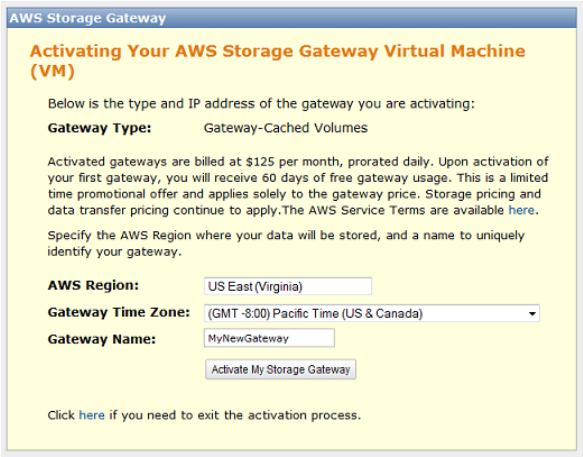
AWS Storage Gateway
Activating Your AWS Storage Gateway Virtual Machine (VM)
Below is the type and IP address of the gateway you are activating:
Gateway Type: Gateway-Cached Volumes
Activated gateways are billed at $125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply.The AWS Service Terms are available here.
Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.
AWS Region: US East (Virginia)
Gateway Time Zone: (GMT -8:00) Pacific Time (US & Canada)
Gateway Name: MyNewGateway
Activate My Storage Gateway
Click here if you need to exit the activation process.

The following screenshot shows the activation page for gateway-VTLs.

AWS Storage Gateway
Activating Your AWS Storage Gateway Virtual Machine
Below is the type of the gateway you are activating:
Gateway Configuration: Gateway-Virtual Tape Library
Activated gateways are billed at $125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply. The AWS Service Terms are available here.
Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.
AWS Region: US West (Oregon)
Gateway Time Zone: (GMT) Western Europe Time, London, Lisbon, Casablanca
Gateway Name: MyGateway-VTL
Medium Changer Type: AWS-Gateway-VTL Learn More
Tape Drive Type: IBM-ULT3580-TD5
Activate My Storage Gateway
Click here if you need to exit the activation process.

iv. Click **Activate My Storage Gateway**.

When the gateway is successfully activated, the AWS Storage Gateway console displays a link to the activated gateway in the **Gateways** section of the **Navigation** pane. To display the entire console, click the gateway you just added. You continue setup by configuring local storage for your gateway. How you configure storage depends on the type of gateway you are setting up.

If activation is not successful, see Troubleshooting Your Gateway-VTL (p. 363) for possible solutions.

### Next Step

You have completed the download, deployment, and activation steps common to all gateways. The remaining setup steps are gateway-specific. To continue the setup process, choose one of the following options:

- Step 2: Configure Local Storage and Alarms (Gateway-Cached and Gateway-Stored) (p. 67)
- Step 2: Configure Local Storage and Alarms (Gateway-VTL) (p. 241)

## Using AWS Storage Gateway with VMware High Availability

VMware High Availability (HA) is a component of vSphere that can provide protection from failures in the infrastructure layer supporting a gateway VM. VMware HA does this by using multiple hosts configured as a cluster so that if a host running a gateway VM fails, the gateway VM can be restarted automatically on another host within the cluster. For more information about VMware HA, go to VMware HA: Concepts and Best Practices.

To use AWS Storage Gateway with VMware HA, we recommend doing the following things:

- Deploy the gateway Open Virtualization Application (OVA) on only one host in a cluster.
- When deploying the OVA, select a data store that is not local to one host. Instead, use a data store that is accessible to all hosts in the cluster. If you select a data store that is local to a host and the host fails, then the data source may not be accessible to other hosts in the cluster and the failover may not succeed.
- Follow the recommended iSCSI settings to prevent your initiator from disconnecting from storage volume targets during failover. In a failover event, it could take between a few seconds to several minutes for a gateway VM to start in a new host in the failover cluster. The recommended iSCSI timeouts for both Windows and Linux clients are greater than the typical time it takes for failover to occur. For more information on customizing Windows clients' timeout settings, see Customizing Your Windows iSCSI Settings (p. 112). For more information on customizing Linux clients' timeout settings, see Customizing Your Linux iSCSI Settings (p. 115).
- With clustering, if you deploy the OVA to the cluster, you will be asked to select a host. Alternately, you can deploy directly to a host in a cluster.

# Hyper-V Host

This section shows you how to perform the tasks described following.

**Topics**


- Step 2.1: Provision a Hyper-V Host to Deploy the AWS Storage Gateway VM (p. 35)
- Step 2.2: Download and Deploy the AWS Storage Gateway VM on Your Host (p. 36)
- Step 2.3: Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 45)
- Step 2.4: Activate Your Storage Gateway (p. 50)


**Next Step**

Step 2.1: Provision a Hyper-V Host to Deploy the AWS Storage Gateway VM (p. 35)

## Step 2.1: Provision a Hyper-V Host to Deploy the AWS Storage Gateway VM

If you have not provisioned a host to deploy the gateway VM, provision a host now.

Setup and Activate Gateway
close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
You will need a host in your datacenter to deploy the gateway virtual machine (VM). Pick a host that meets these minimum requirements.
« Back
Continue

### To provision a host

1. Review the minimum host requirements in Requirements (p. 11).
2. Set up a host in your data center with the Microsoft Hyper-V host. For a minimum set of instructions to install the hypervisor, see Appendix D: Configuring a Microsoft Hyper-V Host for AWS Storage Gateway (p. 411)
3. In the setup wizard, click **Continue**.

### Next Step

Step 2.2: Download and Deploy the AWS Storage Gateway VM on Your Host (p. 36)

## Step 2.2: Download and Deploy the AWS Storage Gateway VM on Your Host

The AWS Storage Gateway virtual machine is available as a Hyper-V downloadable package. This section shows you how to perform the tasks described following.

### Topics

- Step 2.2.1: Download the AWS Storage Gateway VM (p. 36)
- Step 2.2.2: Deploy the AWS Storage Gateway VM to Your Host (p. 37)

### Step 2.2.1: Download the AWS Storage Gateway VM

### To download the VM

1. In the AWS Storage Gateway console's **Setup and Activate Gateway** wizard, navigate to the **DOWNLOAD AND DEPLOY VM** page.
2. Select the version of Microsoft Hyper-V hypervisor you want to run your AWS Storage Gateway VM on, and then click **Continue**. In this example, we select **I want to run the AWS Storage Gateway on Microsoft Hyper-V 2008 R2**.

Setup and Activate Gateway
Close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
Select a virtualization platform to run the AWS Storage Gateway Virtual Machine (VM).
I want to run the AWS Storage Gateway on VMware ESXi
I want to run the AWS Storage Gateway on Microsoft Hyper-V 2008 R2
I want to run the AWS Storage Gateway on Microsoft Hyper-V 2012
« Back
Continue

3. Click **Download** to download a `.zip` file that contains the VM. Extract and save the `.zip` file to a location on your computer, and make note of the location of the folder that was created.

   **Note**
   The `.zip` file is over 500 MB in size and might take some time to download, depending on your network connection.

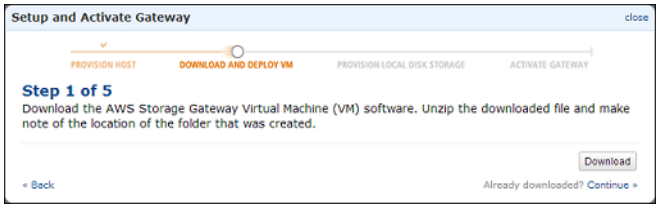
Setup and Activate Gateway
close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
Step 1 of 5
Download the AWS Storage Gateway Virtual Machine (VM) software. Unzip the downloaded file and make note of the location of the folder that was created.
Download
« Back
Already downloaded? Continue »

### Step 2.2.2: Deploy the AWS Storage Gateway VM to Your Host

To work with your hypervisor host, you must connect to it. After you connect to it, you specify locations where VM items are stored, import the VM, and then configure a network for it.

In this exercise, we use Microsoft Hyper-V 2008 R2 Manager to show you how to deploy the VM. Depending on the version of Hyper-V Manager you are using, the screenshots in the exercise might look slightly different from what you see in your UI. However, the procedure is similar for all versions of Microsoft Hyper-V that AWS Storage Gateway supports. For information about the Hyper-V versions that AWS Storage Gateway supports, see Requirements (p. 11).

#### To connect to the hypervisor host

1. Start the Microsoft Hyper-V Manager on your Windows client, and then in the **Actions** pane, click **Connect to Server**.

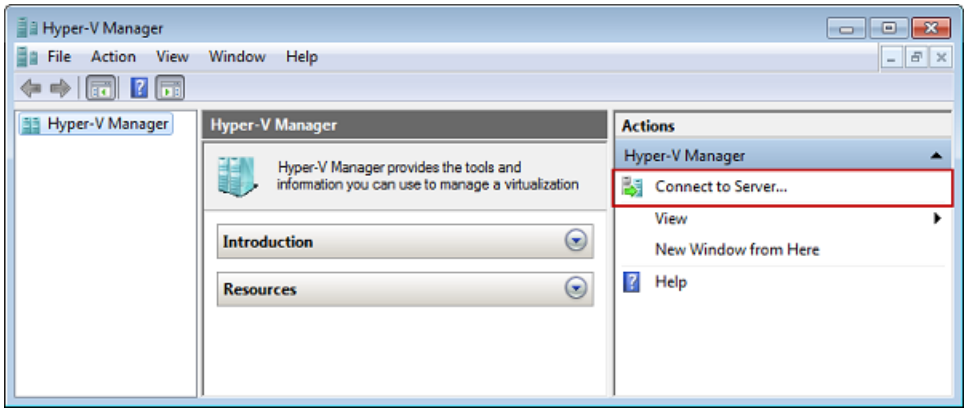


2. In the **Select Computer** dialog box, select **Another Computer**, enter the IP address or host name, and then click **OK**.

   **Note**
   In this exercise, we use `hyperv-server` as a host. Your host name will be different. If you don't find your host name in the **Select Computer** dialog box, you might need to add an entry to your hosts file so that Hyper-V Manager can resolve the server name.

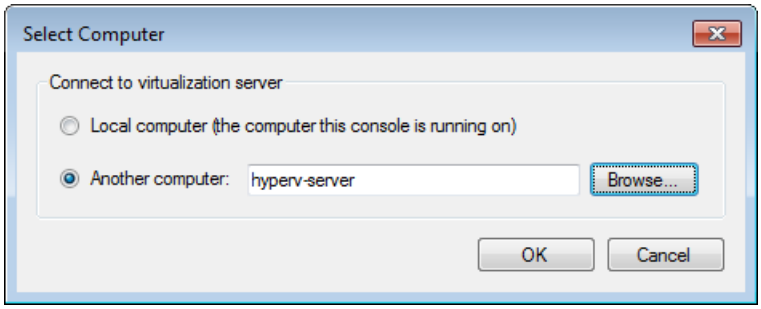


   Your Microsoft Hyper-V Manager is now connected to your host computer.

Now that you are connected to your host, the next step is to create folders on the host to store the downloaded source VM, the imported running VM, and virtual hard disks for the running VM.

#### To specify a location for the virtual hard disks and VM

1. Create locations on the hypervisor host for the gateway virtual hard disks and VM.

**Note**
For Microsoft Hyper-V 2012, this step can be found in the **Choose Folders to Store Virtual Hard Disks** dialog box.

a. Navigate to the hypervisor drive.

   For example, assuming the name of the host in this exercise and the C drive as the correct drive for the host, in the **Start** menu you can type `\\hyperv-server\c$`.

b. Create a folder called `getting-started` with two subfolders, `unzippedSourceVM` and `gateway`.

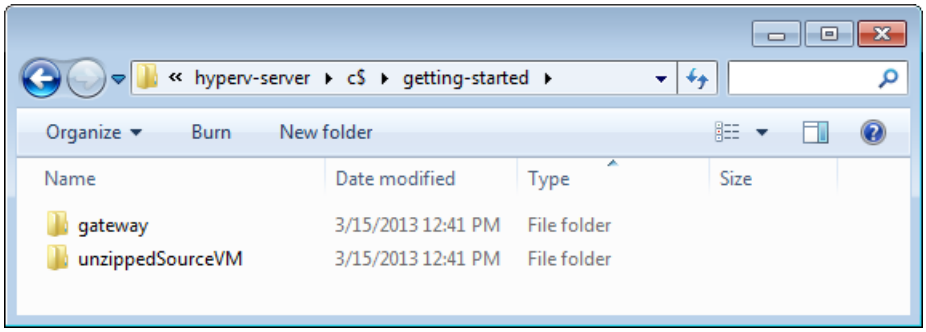


2. Configure the Hyper-V Manager to point to the `gateway` folder you created. The running VM stores its configuration in this folder.

   a. In the **Actions** menu, click **Hyper-V Settings**.

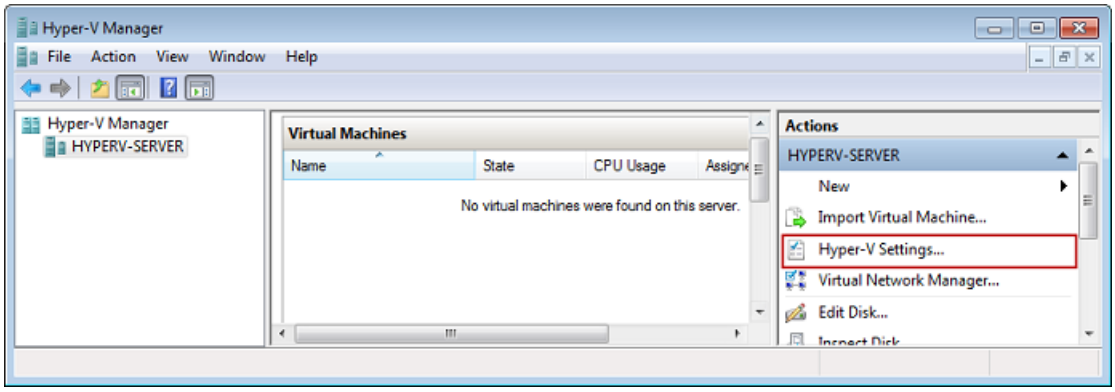


   b. In the **Hyper-V Settings** dialog box, configure the location of the virtual hard disks and virtual machines.

      i. In the left pane, under **Server**, select the **Virtual Hard Disks** setting.

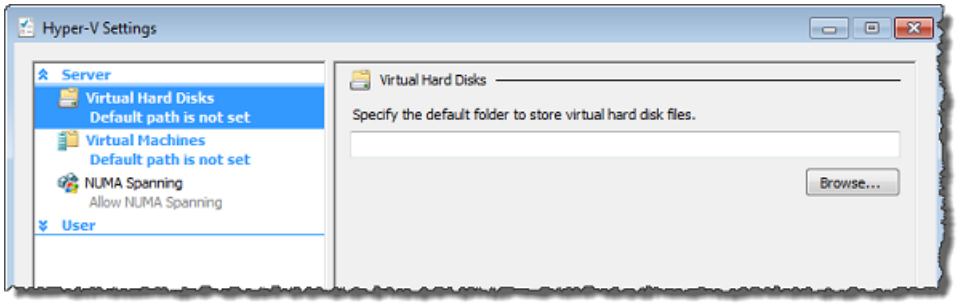

ii. Browse to find the `gateway` folder you created earlier. Note that you are browsing on the hypervisor (host) server.

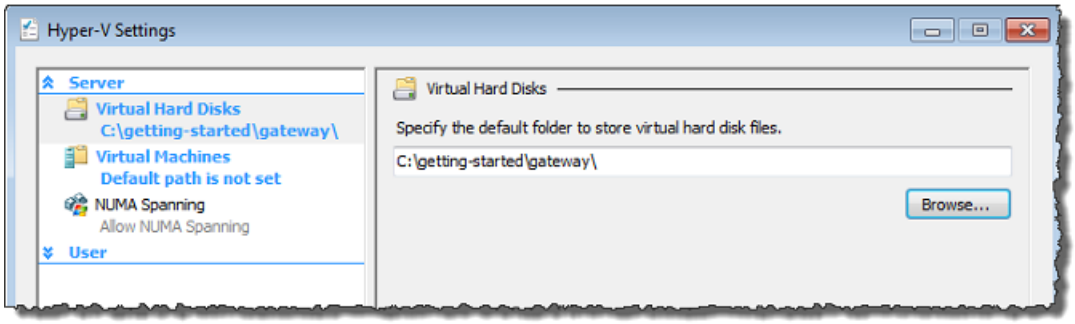


iii. In the left pane, select the **Virtual Machines** setting under **Server**, and then browse to set the location to the same gateway folder and click **OK**.

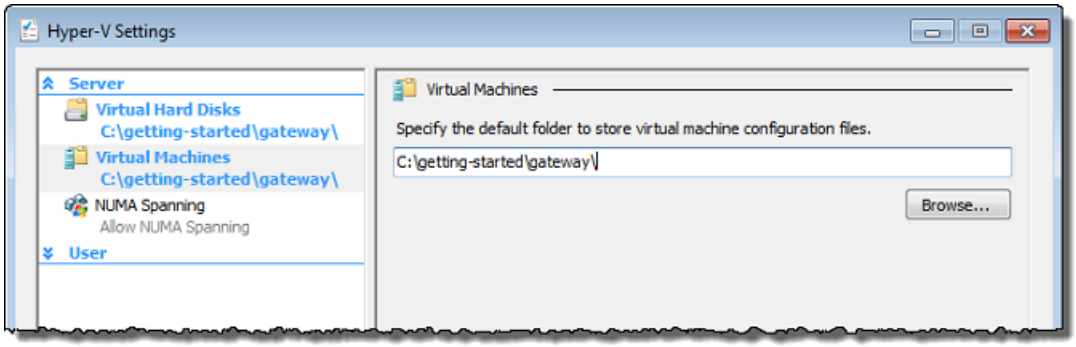


### To import the VM

1. Copy the unzipped source VM files to the folder you created on the host computer. In this exercise, the path is `\\hyperv-server\c$\getting-started\unzippedSourceVM\AWS-Storage-Gateway`.
2. Import the AWS Storage Gateway VM to the host.
   a. In the Hyper-V Manager, select the host **hyperv-server** in the left pane, which shows the console tree.
   b. In the **Actions** menu, click **Import Virtual Machine**.

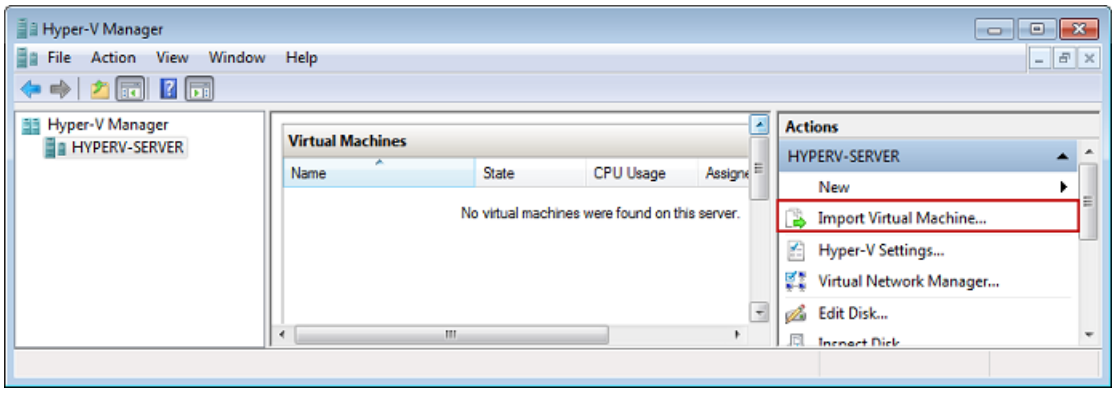


c. In the **Import Virtual Machine** dialog box:

   i. In the **Location** box, find the location you created previously: `\\hyperv-server\c$\getting-started\unzippedSourceVM\AWS-Storage-Gateway`.

      **Note**
      For Microsoft Hyper-V 2012, this step can be found on the **Choose Folders for Virtual Machine Files** dialog box (next dialog box).

      **Caution**
      You must point to the correct folder for the import to succeed. The correct folder to select (`AWS-Storage-Gateway`) contains three other folders (`Snapshots`, `Virtual Hard Disks`, `Virtual Machines`) and one file (`config.xml`). Depending on how you unzip the gateway source files, you might end up with an extra folder level. For help troubleshooting imports, see Troubleshooting Your Microsoft Hyper-V Setup (p. 417).

   ii. Select **Copy the virtual machine (create a new unique ID)**.

      **Note**
      For Microsoft Hyper-V 2012, this step can be found in the **Choose Import Type** tab.

   iii. Check **Duplicate all files so the same virtual machine can be imported again**.

   iv. Click **Import**.

      **Caution**
      It is important to select the **Copy the virtual machine (create a new unique ID)** and **Duplicate all files so the same virtual machine can be imported again** options, especially if you intend to reuse the unzipped gateway source files.

      **Important**
      You must have 75 GiB of disk space for installation of the VM image and system data.

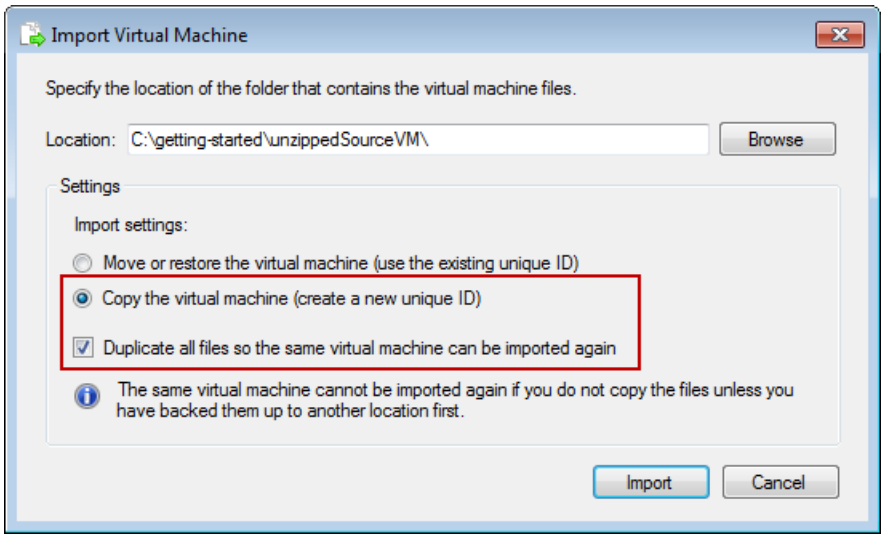
Import Virtual Machine
Specify the location of the folder that contains the virtual machine files.
Location:
C:\getting-started\unzippedSourceVM\
Browse
Settings
Import settings:
Move or restore the virtual machine (use the existing unique ID)
Copy the virtual machine (create a new unique ID)
Duplicate all files so the same virtual machine can be imported again
The same virtual machine cannot be imported again if you do not copy the files unless you have backed them up to another location first.
Import
Cancel

After the import is complete, a virtual machine named AWS-Storage-Gateway is created.

3. Rename the virtual machine to avoid confusion with other virtual machines that you might import to the host.

   a. Right-click the virtual machine name, and then click **Rename**.
   b. Provide a new name for the virtual machine.

      In this exercise, we will use `ExampleGatewayHyperV`.

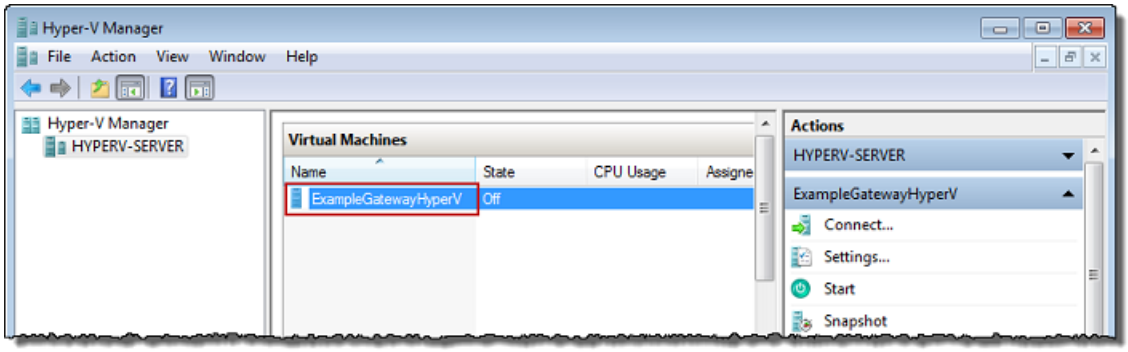


4. Confirm that **Time synchronization** for the VM is selected in **Integration Services**.

   a. In the **Virtual Machines** list, select the virtual machine **ExampleGatewayHyperV**.
   b. In the **Actions** menu, click **Settings**.
   c. In the **Settings** dialog box, select **Integration services** under **Management** and confirm that **Time synchronization** is checked.

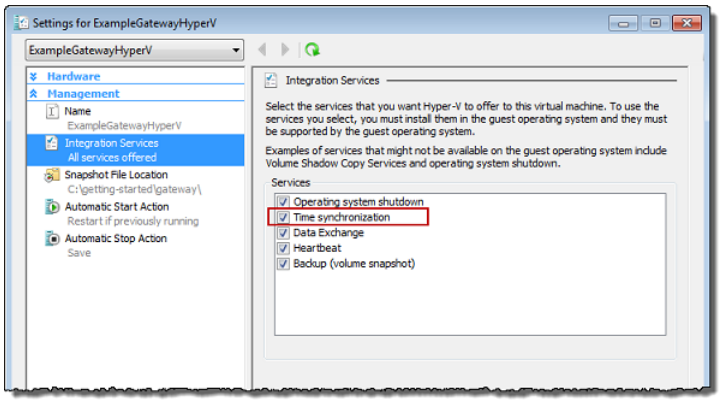
Settings for ExampleGatewayHyperV
ExampleGatewayHyperV
Hardware
Management
Name
ExampleGatewayHyperV
Integration Services
All services offered
Snapshot File Location
C:\getting-started\gateway\
Automatic Start Action
Restart if previously running
Automatic Stop Action
Save
Integration Services
Select the services that you want Hyper-V to offer to this virtual machine. To use the services you select, you must install them in the guest operating system and they must be supported by the guest operating system.
Examples of services that might not be available on the guest operating system include Volume Shadow Copy Services and operating system shutdown.
Services
Operating system shutdown
Time synchronization
Data Exchange
Heartbeat
Backup (volume snapshot)

5. Configure the host time if you have not already done so, and then click **OK**.

   It is important to make sure that your host clock is set to the correct time. The following steps show you how to set the time by using the Server Configuration Tool (`Sconfig.cmd`). For more information about `Sconfig.cmd`, go to Configure a Server Core Server with Sconfig.cmd. (Depending on the version of Microsoft Hyper-V you are running, you might be able to set the time in other ways.)

   a. Access the `Sconfig.cmd` tool by either using the hypervisor host console or logging in remotely.

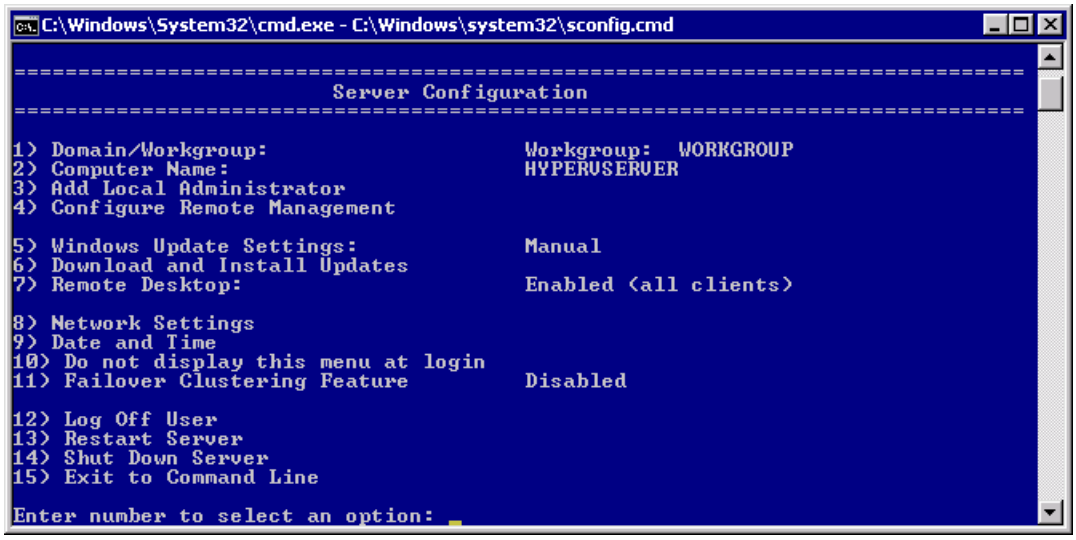
C:\Windows\System32\cmd.exe - C:\Windows\system32\sconfig.cmd
Server Configuration
1) Domain/Workgroup: Workgroup: WORKGROUP
2) Computer Name: HYPERVSERVER
3) Add Local Administrator
4) Configure Remote Management
5) Windows Update Settings: Manual
6) Download and Install Updates
7) Remote Desktop: Enabled (all clients)
8) Network Settings
9) Date and Time
10) Do not display this menu at login
11) Failover Clustering Feature Disabled
12) Log Off User
13) Restart Server
14) Shut Down Server
15) Exit to Command Line
Enter number to select an option:

   b. Enter option 9 **Date and Time**.

      The **Date and Time** control panel appears.

   c. Configure the time, and then click **OK**.

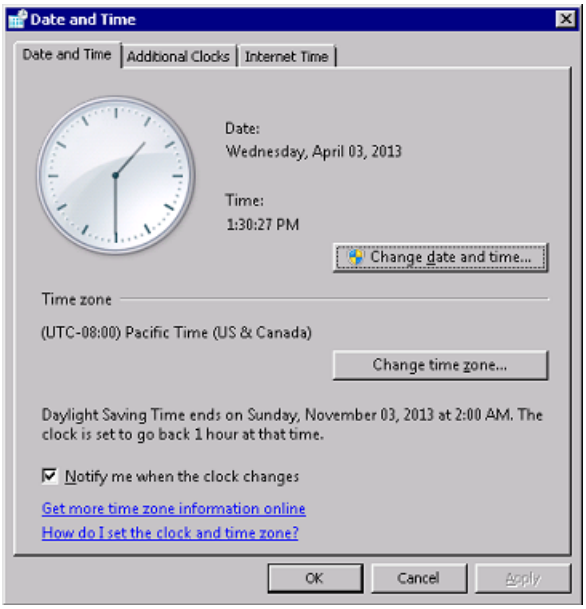
Date and Time
Date and Time
Additional Clocks
Internet Time
Date:
Wednesday, April 03, 2013
Time:
1:30:27 PM
Change date and time...
Time zone
(UTC-08:00) Pacific Time (US & Canada)
Change time zone...
Daylight Saving Time ends on Sunday, November 03, 2013 at 2:00 AM. The clock is set to go back 1 hour at that time.
Notify me when the clock changes
Get more time zone information online
How do I set the clock and time zone?
OK
Cancel
Apply

Now you can log off the remote computer.

### To configure a virtual network and use it for the VM

1. Configure virtual network settings for the Hyper-V host.

   **Note**
   In this exercise, we assume the host does not have virtual network settings configured. If you already have a virtual network configured, go to step 2.

   a. In the **Actions** menu, under the hypervisor host name (for example, `hyperv-server`), click **Virtual Network Manager**.

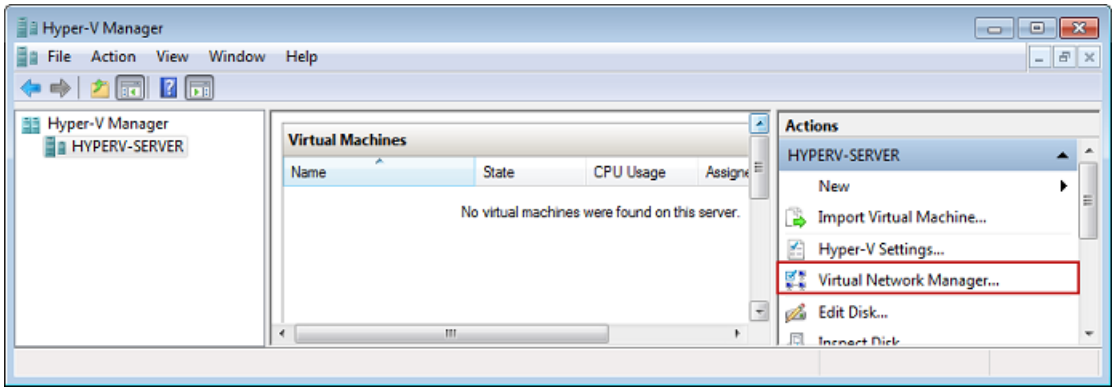


   b. In the **Virtual Network Manager** dialog box, select **New virtual network**.

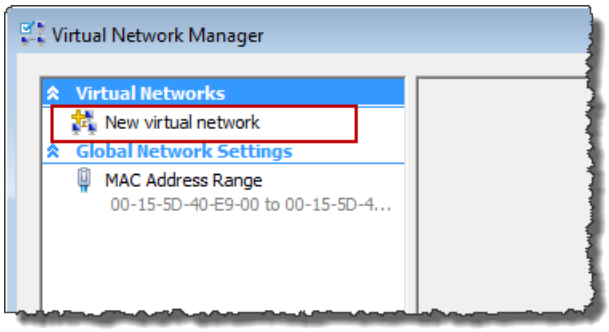


c. Select **External** as the virtual network type, and then click **Add**.

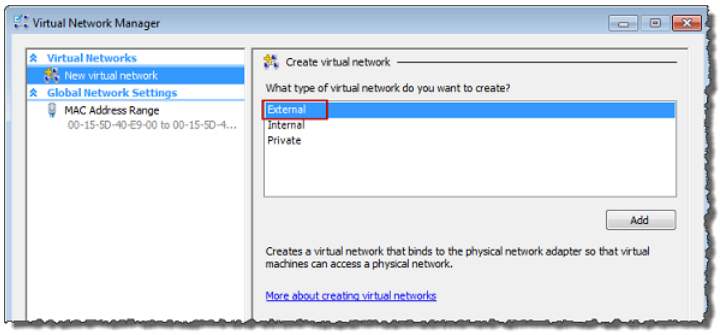
Virtual Network Manager
Virtual Networks
New virtual network
Global Network Settings
MAC Address Range
00-15-5D-40-E9-00 to 00-15-5D-4...
Create virtual network
What type of virtual network do you want to create?
External
Internal
Private
Add
Creates a virtual network that binds to the physical network adapter so that virtual machines can access a physical network.
More about creating virtual networks

d. Provide a name for the network, and then click **OK**.

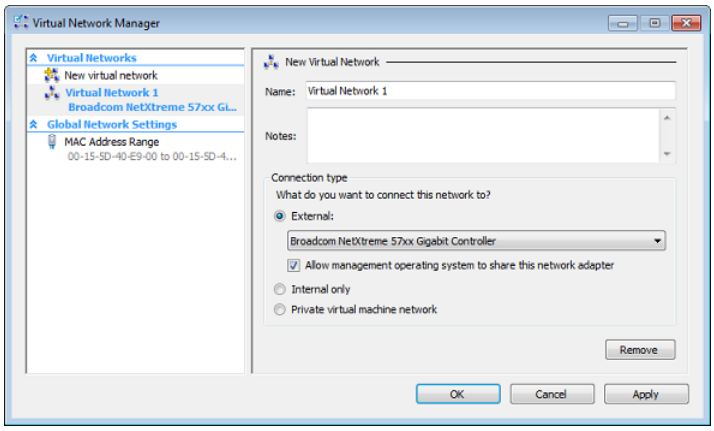


2. Configure the virtual machine to use the virtual network.

   a. In the **Virtual Machines** list, select the virtual machine **ExampleGatewayHyperV**.
   b. In the **Actions** pane, select **Settings**.

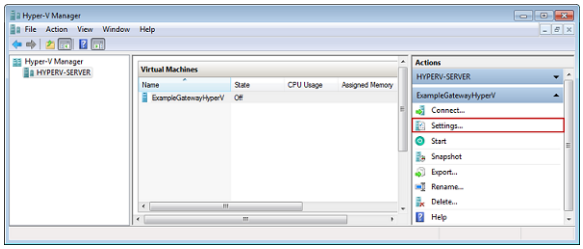

c. In the **Settings** window, select **Network Adapter**. The **Network Adapter** should have a status of **Not connected**.

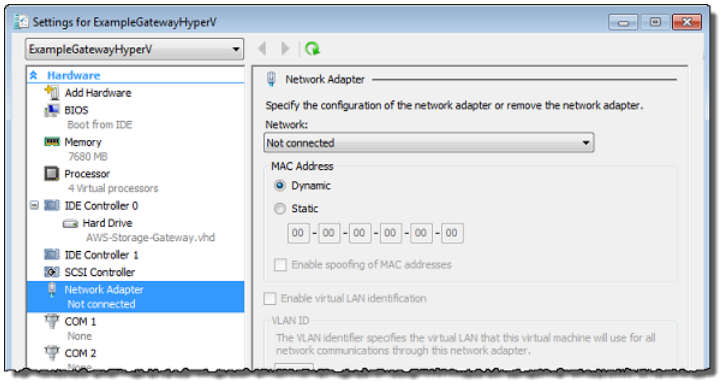


d. In the **Network** box at right, select a network, and then click **OK**. In the following screenshot, **Virtual Network 1** is selected.

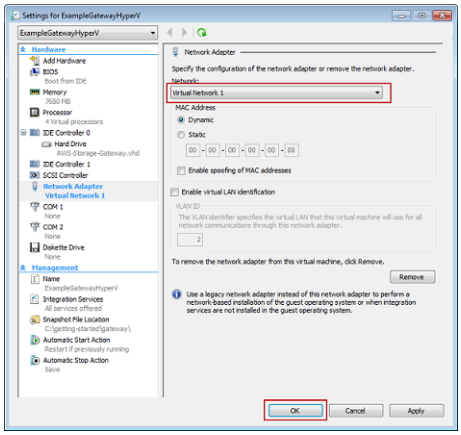

**Next Step**

Step 2.3: Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 45)

## Step 2.3: Provision Local Storage for the AWS Storage Gateway VM (Hyper-V)

Next, you allocate local storage for your deployed gateway VM. After you activate your gateway, you configure this local storage for the gateway's use.

**Topics**


- Step 2.3.1: Decide the Amount of Local Disk Storage (p. 45)
- Step 2.3.2: Provision Local Disk Storage (p. 46)


### Step 2.3.1: Decide the Amount of Local Disk Storage

In this step, you decide the number and size of disks to allocate for your gateway, as follows:

- All gateways require one or more disks for an upload buffer.

- Depending on the storage solution you deploy (see Plan Your AWS Storage Gateway Deployment (p. 10)), the gateway requires the following additional storage:
  - Gateway-VTLs and gateway-cached volumes require one or more disks for cache storage.
  - Gateway-stored volumes require one or more disks for volume storage—one local disk for each volume created on the gateway.

You will need at least two disks to begin with. The following table recommends sizes for the local disk storage you will allocate in this step.

| Disk Size | Gateway-Cached | Gateway-Stored | Gateway-VTL |
|---|---|---|---|
| Upload buffer (minimum) | 150 GiB | 150 GiB | 150 GiB |
| Upload buffer (maximum) | 2 TiB | 2 TiB | 2 TiB |
| Cache storage (minimum) | 1.1 x upload buffer | — | 1.1 x upload buffer |
| Cache storage (maximum) | 16 TiB | — | 16 TiB |
| Volume storage (minimum) | — | 1GiB | — |
| Volume storage (maximum) | 32 TiB | 1 TiB | — |

Before going to the next step, decide the number and size of disks to allocate. You can add more local storage after you set the gateway up, as your workload demands. If you're not sure what number and size of disks to use, provision two local disks with minimum recommended size.

If you plan to deploy your gateway in production, you should consider your real workload in determining disk sizes. For disk size guidelines, see one of the following sections depending on the type of gateway you are deploying:

- Gateway-cached: Managing the Upload Buffer and Cache Storage (Gateway-Cached) (p. 169)
- Gateway-stored: Managing the Upload Buffer (Gateway-Stored) (p. 178)
- Gateway-VTL: Managing the Upload Buffer and Cache Storage (Gateway-VTL) (p. 304)

For more information about how gateways use local storage, see How AWS Storage Gateway Works (Architecture) (p. 2).

In the next step, you allocate the local disk storage to the gateway VM you deployed.

### Step 2.3.2: Provision Local Disk Storage

You can add disks from direct attached storage (DAS) or from a storage area network (SAN). This section tells how to add a virtual disk from a DAS. For instructions on attaching iSCSI volumes from an existing SAN for this step, see To add a new iSCSI target (p. 407).

Repeat the following procedure for each local disk you planned in the preceding step.

#### To allocate a local disk from DAS

1. Start the Microsoft Hyper-V Manager, and then connect to the hypervisor.
2. In the **Virtual Machines** list, select the virtual machine **ExampleGatewayHyperV**.
3. In the **Actions** pane of your gateway, select **Settings**.

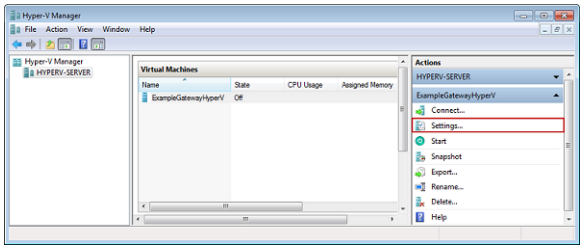


4. In the **Settings** window, select **SCSI Controller**, and then click **Add**.

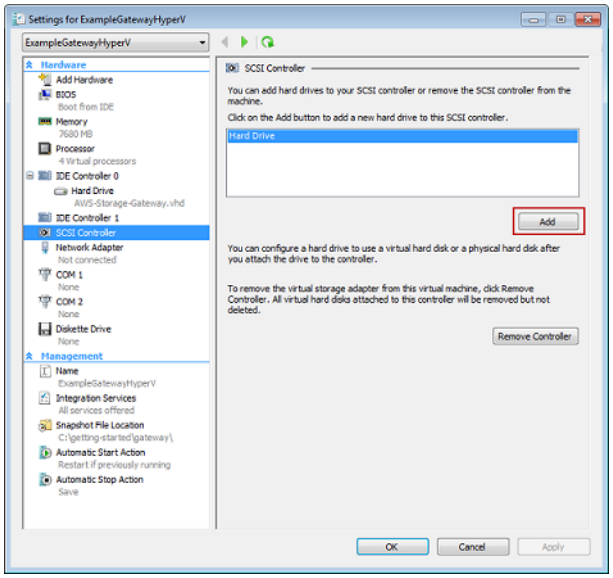


5. In the **Hard Drive** pane, under **Media**, click **New**.

   **Note**

   In this example, notice that **virtual hard disk (.vhd) file** is selected. AWS Storage Gateway supports .vhdx file type. This file type enables you to create larger virtual disks than other file types. If you create a .vhdx type virtual disk, make sure that the size of the virtual disks you create does not exceed the recommended disk size for your gateway. For more information, see Step 2.3.1: Decide the Amount of Local Disk Storage (p. 45).

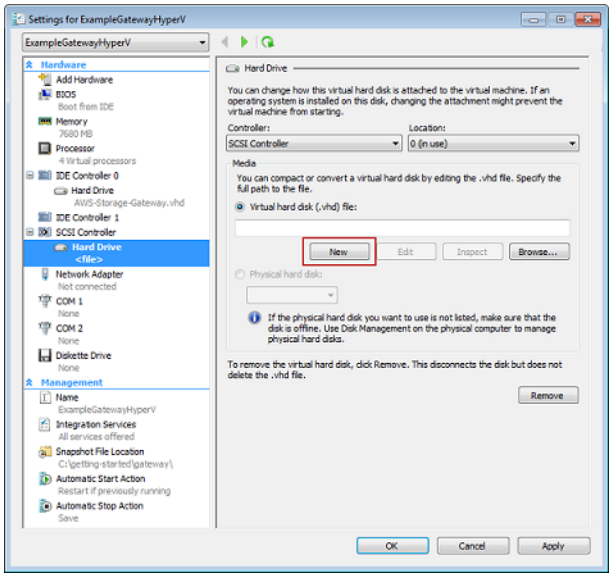

6. In the **New Virtual Hard Disk Wizard**, create a new virtual hard disk.

   a. On the **Before You Begin** page, click **Next**.

      **Note**
      If you are using the Microsoft Hyper-V 2012 Hypervisor, you will be prompted to choose a disk format (**VHD** or **VHDX**).

   b. On the **Choose Disk Type** page, choose **Fixed size**, and then click **Next**.

      When you use fixed-size provisioning, the disk storage is allocated immediately, resulting in better performance. If you don't use fixed-size provisioning, the storage is allocated on demand. On-demand allocation can affect the functioning of AWS Storage Gateway. For AWS Storage Gateway to function properly, the VM disks must be stored in fixed-size provisioned format.

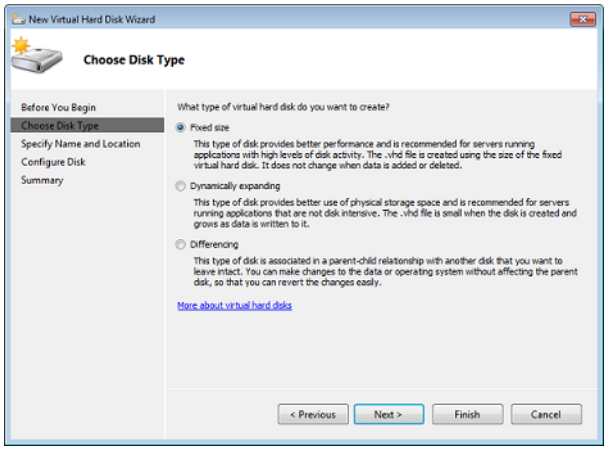
New Virtual Hard Disk Wizard
Choose Disk Type
Before You Begin
Choose Disk Type
Specify Name and Location
Configure Disk
Summary
What type of virtual hard disk do you want to create?
Fixed size
This type of disk provides better performance and is recommended for servers running applications with high levels of disk activity. The .vhd file is created using the size of the fixed virtual hard disk. It does not change when data is added or deleted.
Dynamically expanding
This type of disk provides better use of physical storage space and is recommended for servers running applications that are not disk intensive. The .vhd file is small when the disk is created and grows as data is written to it.
Differencing
This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily.
More about virtual hard disks
< Previous
Next >
Finish
Cancel

   c. On the **Specify Name and Location** page, specify a name and location for the virtual hard disk.

      i. Specify `UploadBuffer.vhd` as the name.

      ii. Specify the location as `c$\getting-started\gateway`.

         **Note**
         In this example setup, you store the virtual disk with the virtual machine. For real-world workloads, we strongly recommend that you not provision local disks using the same underlying physical storage disk. Depending on your hosting environment, performance, and portability requirements, consider selecting a different physical disk in this step.

      iii. Click **Next**.

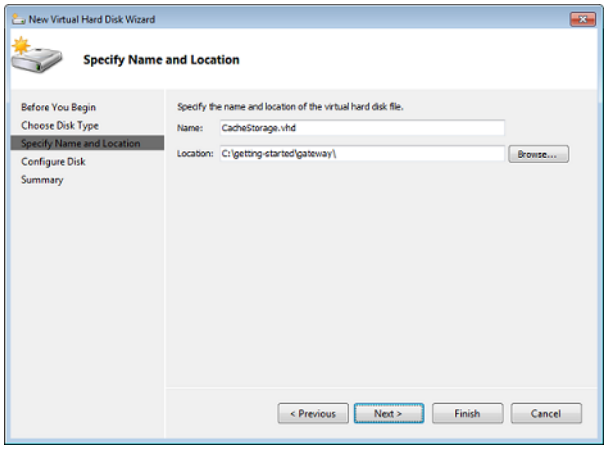

d. In the **Configure Disk** page, specify the disk size you previously decided, and click **Finish**.

> **Note**
> The disk sizes used in these examples in this documentation are not suitable for real-world workloads. If your setup is a trial setup, you can use small disk sizes.

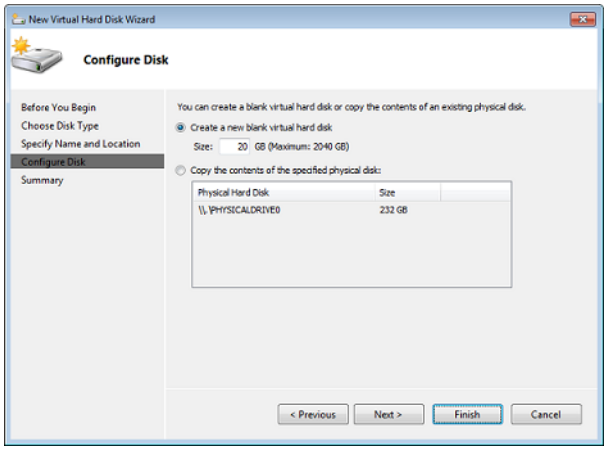


e. After the virtual disk is created, verify that **Hard Drive** shows up under **SCSI controller**.

f. Click **SCSI Controller** to prepare to add another hard drive, and then click **OK**.

> **Warning**
> When you add an additional hard drive, you must first click **SCSI Controller** and then follow the steps in the following procedure. Clicking **New** when viewing the details of an existing hard drive will replace the existing drive.

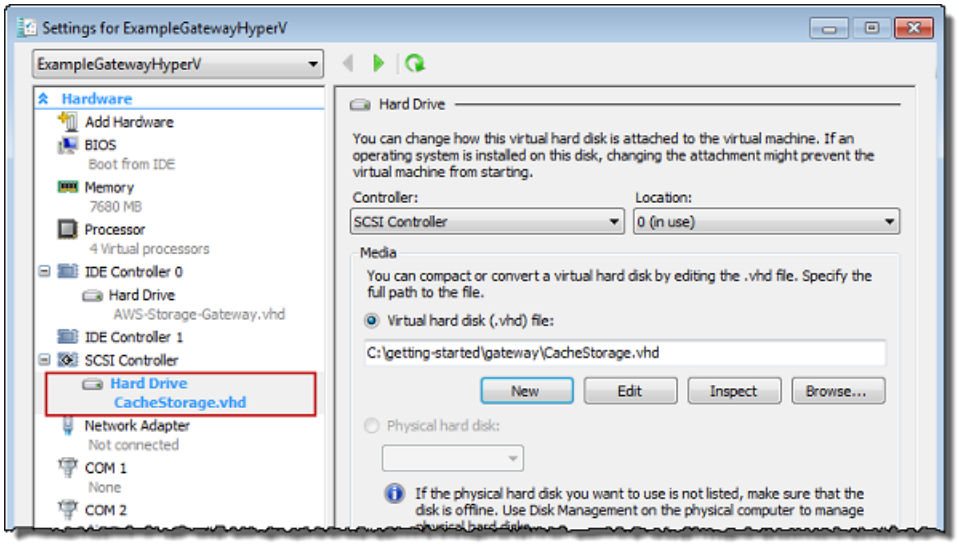


7. Repeat steps 1 through 6 for each disk to allocate.
8. Verify your gateway VM has the disks you added.

   a. Start the Microsoft Hyper-V Manager and connect to the hypervisor if it is not already started.

   b. In the **Virtual Machines** list, select the virtual machine **ExampleGatewayHyperV**.

   c. In the **Actions** pane, select **Settings,** and then in the **Hardware** pane select **SCSI Controller** and verify there are two disks under the SCSI controller.

The two disks you created are used later in the AWS Storage Gateway console and appear as **SCSI (0:0)** and **SCSI (0:1)** in drop-down lists. In the following example, the **CacheStorage.vhd** disk is selected and is SCSI (0:0).

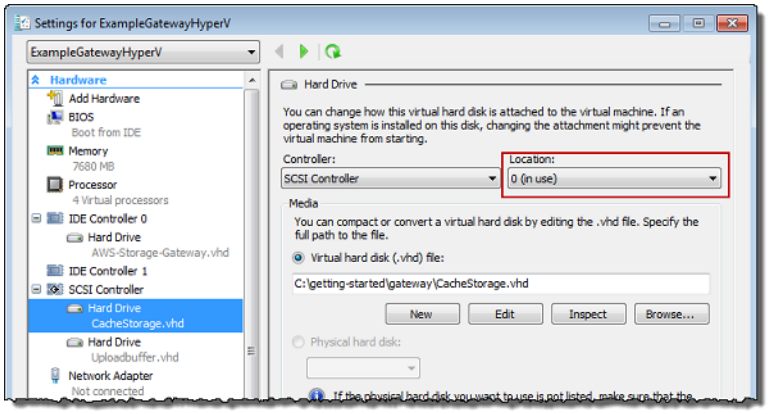


**Next Step**

Step 2.4: Activate Your Storage Gateway (p. 50)

## Step 2.4: Activate Your Storage Gateway

Now, you can activate your gateway. The activation process associates your gateway with your AWS account. You must turn on the gateway VM before you activate your gateway.

### To activate your gateway

1. Turn on the VM.

   a. Start the Microsoft Hyper-V Manager and connect to the hypervisor if it is not already connected.
   b. In the **Virtual Machines** list, select the virtual machine you deployed, **ExampleGatewayHyperV**.
   c. In the **Actions** pane for your gateway, click **Start**.

      The **Virtual Machine Connection** window appears. If this window does not appear, click **Connect** in the **Actions** pane for your gateway.

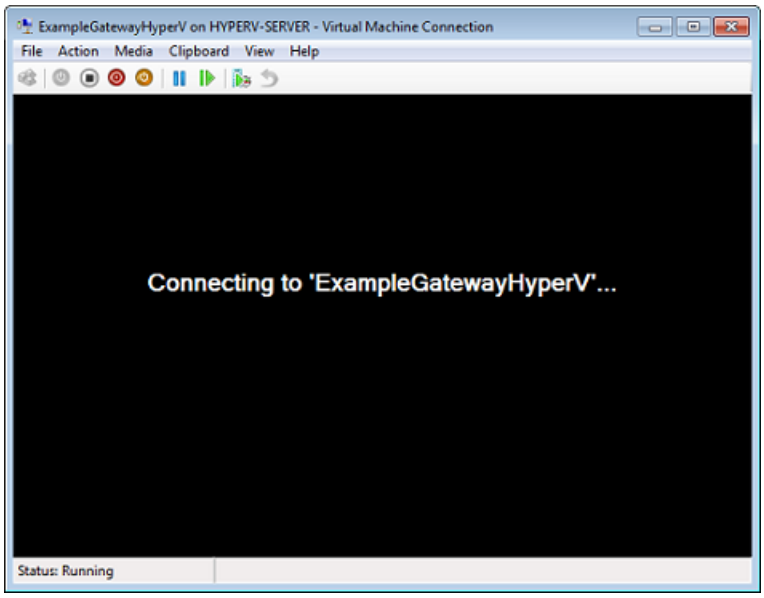
ExampleGatewayHyperV on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
Connecting to 'ExampleGatewayHyperV'...
Status: Running

d. If an authentication window appears, type the user name and password provided to you by the hypervisor administrator.

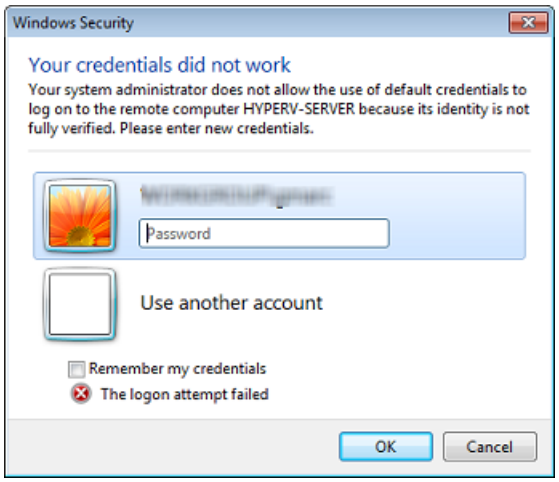
Windows Security
Your credentials did not work
Your system administrator does not allow the use of default credentials to log on to the remote computer HYPERV-SERVER because its identity is not fully verified. Please enter new credentials.
Password
Use another account
Remember my credentials
The logon attempt failed
OK
Cancel

e. After a few moments, the virtual machine is ready for you to log in.

The example following shows the login prompt you see when the VM is ready.

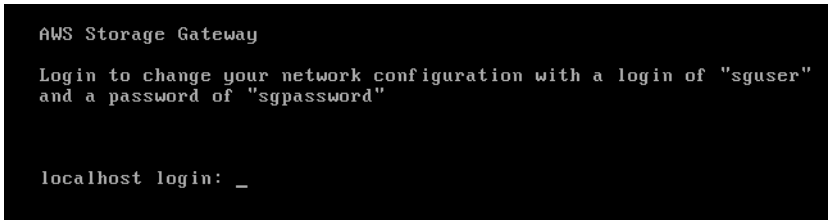
AWS Storage Gateway
Login to change your network configuration with a login of "sguser"
and a password of "sgpassword"
localhost login: _

2. Activate your gateway.

   a. Obtain the IP address of your gateway.

i. In Microsoft Hyper-V Manager, select the deployed gateway VM.
ii. In the **Virtual Machines** list, select the virtual machine **ExampleGatewayHyperV**.
iii. In the **Actions** pane, click **Connect**.

The **Virtual Machine Connection** window appears.

iv. At the login prompt, type the user name `sguser` and the password `sgpassword`.
v. In the **AWS Storage Gateway Configuration** menu, select option 2, **Static IP Address Configuration**.

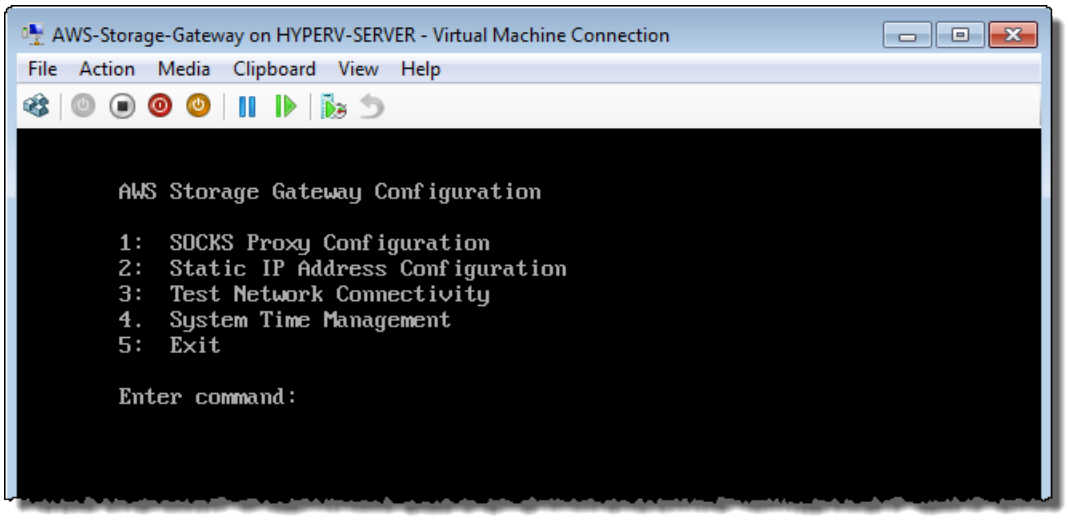
AWS-Storage-Gateway on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
AWS Storage Gateway Configuration
1: SOCKS Proxy Configuration
2: Static IP Address Configuration
3: Test Network Connectivity
4. System Time Management
5: Exit
Enter command:

vi. In the **AWS Storage Gateway Static IP Address Configuration** menu, select option 1, **View Network Configuration**.

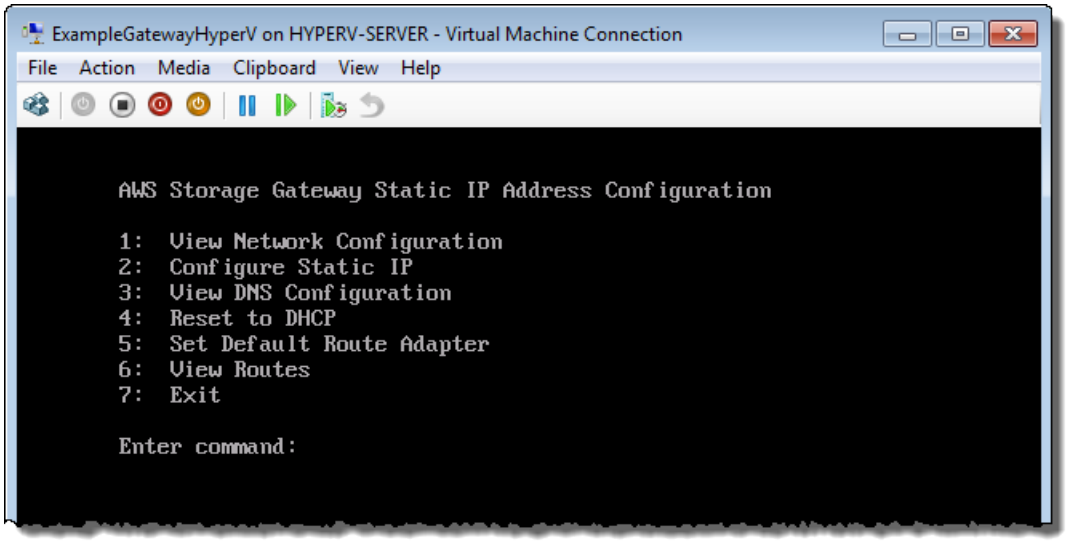
ExampleGatewayHyperV on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
AWS Storage Gateway Static IP Address Configuration
1: View Network Configuration
2: Configure Static IP
3: View DNS Configuration
4: Reset to DHCP
5: Set Default Route Adapter
6: View Routes
7: Exit
Enter command:

vii. Type the identifier of the adapter.

In most setups, `eth0` will be the adapter identifier.

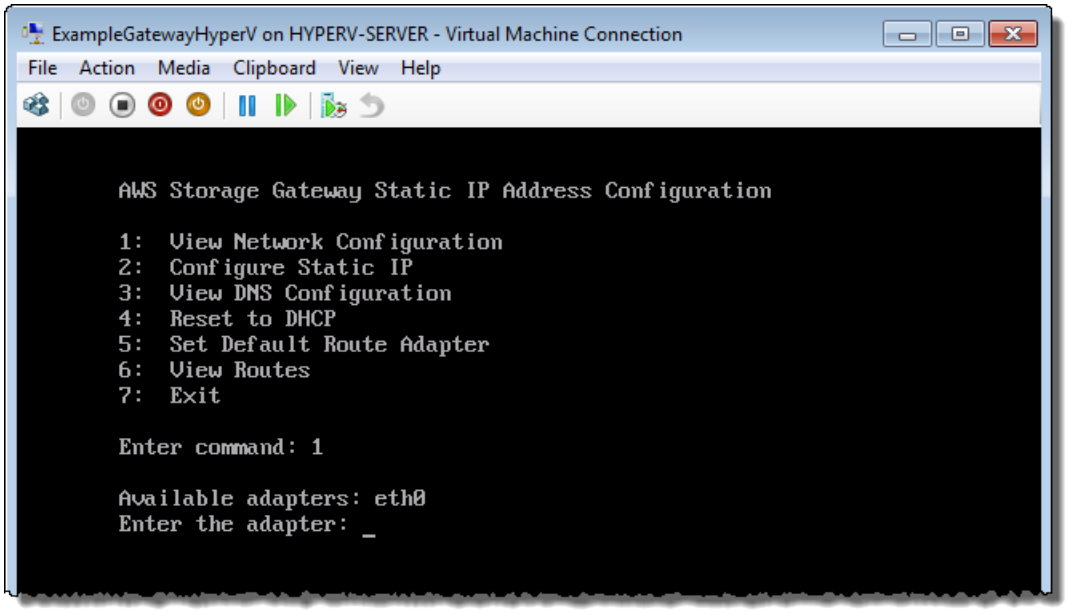
ExampleGatewayHyperV on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
AWS Storage Gateway Static IP Address Configuration
1: View Network Configuration
2: Configure Static IP
3: View DNS Configuration
4: Reset to DHCP
5: Set Default Route Adapter
6: View Routes
7: Exit
Enter command: 1
Available adapters: eth0
Enter the adapter: _

viii. Get the IP address from the adapter information.

In the example below, the IP address is 10.61.64.130. Your gateway's IP address will be different.

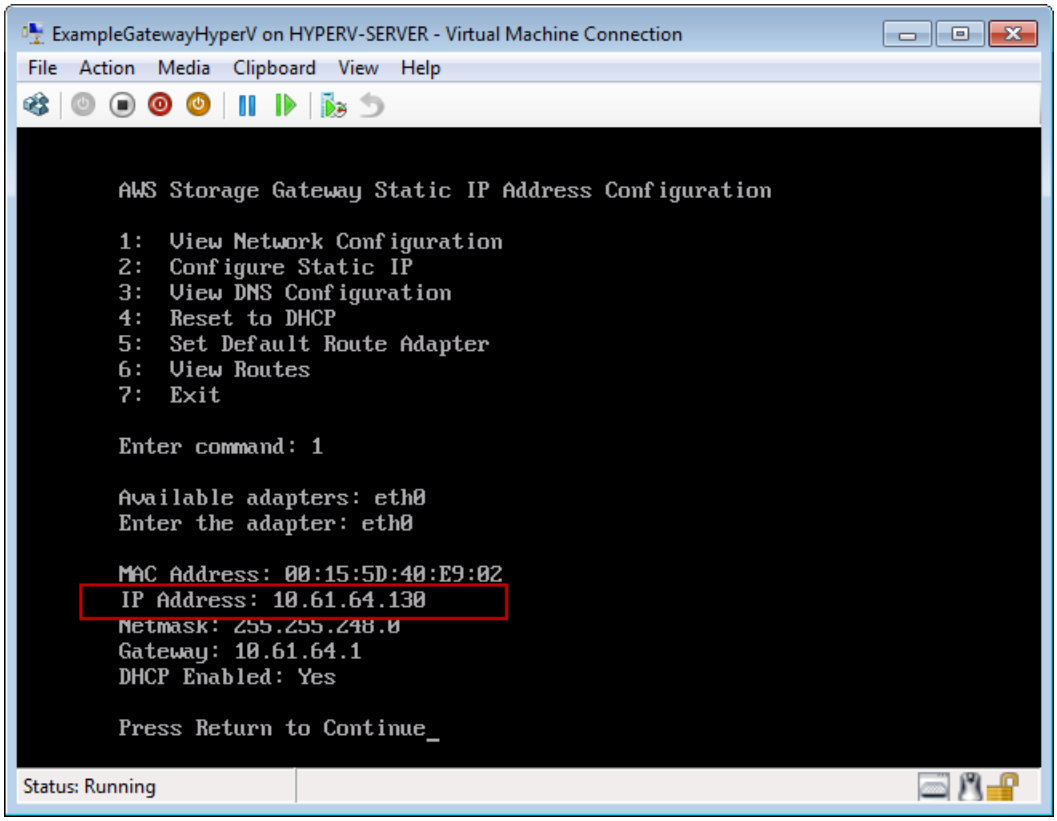


ix. Press the Enter key, and follow the prompts to exit the configuration menu.

b. Associate your gateway with your AWS account.

i. In the AWS Storage Gateway console's **Setup and Activate Gateway** wizard, navigate to the **ACTIVATE GATEWAY** page.

If the wizard has not already started, in the navigation pane, click the **Deploy a New Gateway** link then click **Continue** in each wizard step in the **Setup and Activate Gateway** until you reach that page.

ii. Type the IP address of your gateway, and click **Proceed to Activation**. Make a note of the IP address, because you will need it later when you access your devices.

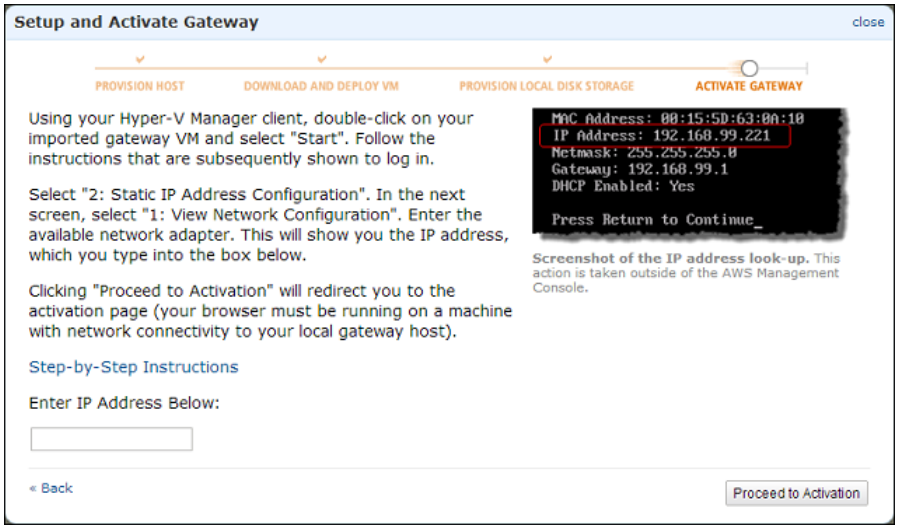
Setup and Activate Gateway
close
PROVISION HOST
DOWNLOAD AND DEPLOY VM
PROVISION LOCAL DISK STORAGE
ACTIVATE GATEWAY
Using your Hyper-V Manager client, double-click on your imported gateway VM and select "Start". Follow the instructions that are subsequently shown to log in.
Select "2: Static IP Address Configuration". In the next screen, select "1: View Network Configuration". Enter the available network adapter. This will show you the IP address, which you type into the box below.
Clicking "Proceed to Activation" will redirect you to the activation page (your browser must be running on a machine with network connectivity to your local gateway host).
Step-by-Step Instructions
Enter IP Address Below:
« Back
Proceed to Activation
MAC Address: 08:15:5D:63:0A:10
IP Address: 192.168.99.221
Netmask: 255.255.255.0
Gateway: 192.168.99.1
DHCP Enabled: Yes
Press Return to Continue_
Screenshot of the IP address look-up. This action is taken outside of the AWS Management Console.

**Note**
During activation, your browser connects to the gateway. If activation fails, check that the IP address you entered is correct. If the IP address is correct, confirm that your network is configured to allow your browser to access the gateway VM.

iii. To complete the activation process, enter the information listed following on the activation page. The activation page you see depends on your gateway configuration.

- **AWS Region** determines where AWS stores your snapshots (for gateway-cached volumes) or virtual tapes (for gateway-VTLs). You select the region in which to deploy your gateway in the **AWS Region** box. You cannot change the region after the gateway is activated.
- **Gateway Time Zone** specifies the time zone to use for your gateway.
- **Gateway Name** identifies your gateway. You use this name to manage your gateway in the console; you can change it after the gateway is activated. This name must be unique to your account.
- **Medium Changer Type** specifies the type of medium changer to use for your backup software (available for gateway-VTLs only).

**Important**
The type of medium changer you select depends on the backup software you plan to use. The following table shows which medium changer to select for your backup software.

| Backup Software | Medium Changer Type |
|---|---|
| Symantec NetBackup Version 7.x | `AWS-Gateway-VTL` |
| Symantec Backup Exec 2012 | `STK-L700` |

| Backup Software | Medium Changer Type |
|---|---|
| Symantec Backup Exec 2014 | `AWS-Gateway-VTL` |
| Symantec Backup Exec 15 | `AWS-Gateway-VTL` |
| Microsoft System Center 2012 R2 Data Protection Manager | `STK-L700` |
| Veeam Backup & Replication V7 | `STK-L700` |
| Veeam Backup & Replication V8 | `STK-L700` |
| Dell NetVault Backup 10.0 | `STK-L700` |

**Note**
You must select the medium changer that is recommended for your backup software. Other medium changers might not function properly. You can select a different medium changer after the gateway is activated. For more information, see Selecting a Medium Changer After Gateway Activation (p. 424).

- **Tape Drive Type** specifies the type of tape drive used by this gateway (available for gateway-VTLs only).

**Note**
AWS Storage Gateway is supported in the following regions.

| Region Name | Region String | Gateway-Cached | Gateway-Stored | Gateway-VTL |
|---|---|---|---|---|
| US East (N. Virginia) | `us-east-1` | yes | yes | yes |
| US West (N. California) | `us-west-1` | yes | yes | yes |
| US West (Oregon) | `us-west-2` | yes | yes | yes |
| EU (Ireland) | `eu-west-1` | yes | yes | yes |
| EU (Frankfurt) | `eu-central-1` | yes | yes | yes |
| Asia Pacific (Singapore) | `ap-south-east-1` | yes | yes | no |
| Asia Pacific (Sydney) | `ap-south-east-2` | yes | yes | yes |
| Asia Pacific (Tokyo) | `ap-north-east-1` | yes | yes | yes |
| South America (Sao Paulo) | `sa-east-1` | yes | yes | no |

The following screenshot shows the activation page for gateway-cached volumes or gateway-stored volumes.

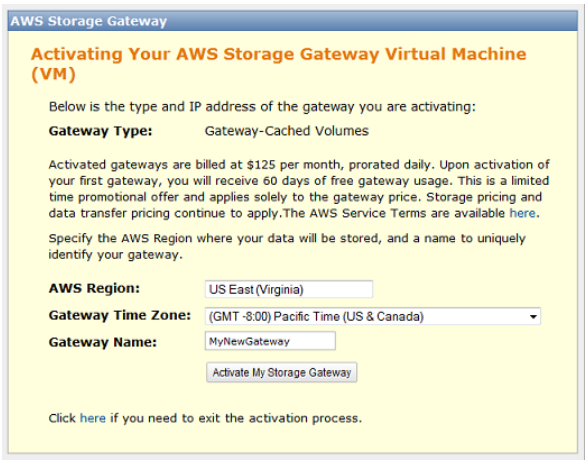
AWS Storage Gateway
Activating Your AWS Storage Gateway Virtual Machine (VM)
Below is the type and IP address of the gateway you are activating:
Gateway Type: Gateway-Cached Volumes
Activated gateways are billed at $125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply.The AWS Service Terms are available here.
Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.
AWS Region: US East (Virginia)
Gateway Time Zone: (GMT -8:00) Pacific Time (US & Canada)
Gateway Name: MyNewGateway
Activate My Storage Gateway
Click here if you need to exit the activation process.

The following screenshot shows the activation page for gateway-VTLs.

AWS Storage Gateway
Activating Your AWS Storage Gateway Virtual Machine
Below is the type of the gateway you are activating:
Gateway Configuration: Gateway-Virtual Tape Library
Activated gateways are billed at $125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply. The AWS Service Terms are available here.
Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.
AWS Region: US West (Oregon)
Gateway Time Zone: (GMT) Western Europe Time, London, Lisbon, Casablanca
Gateway Name: MyGateway-VTL
Medium Changer Type: AWS-Gateway-VTL Learn More
Tape Drive Type: IBM-ULT3580-TD5
Activate My Storage Gateway
Click here if you need to exit the activation process.

iv. Click **Activate My Storage Gateway**.

When the gateway is successfully activated, the AWS Storage Gateway console displays a link to the activated gateway in the **Gateways** section of the **Navigation** pane. To display the entire console, click the gateway you just added. You continue setup by configuring local storage for your gateway. How you configure storage depends on the type of gateway you are setting up.

If activation is not successful, see Troubleshooting Your Gateway-VTL (p. 363) for possible solutions.

### Next Step

You have completed the download, deployment, and activation steps common to all gateways. The remaining setup steps are gateway-specific. To continue the setup process, choose one of the following options:

- Step 2: Configure Local Storage and Alarms (Gateway-Cached and Gateway-Stored) (p. 67)

- Step 2: Configure Local Storage and Alarms (Gateway-VTL) (p. 241)

# Amazon EC2 Gateway

This section shows you how to deploy and activate a gateway on an Amazon EC2 instance. Gateways hosted on Amazon EC2 instances can be used for disaster recovery and data mirroring. You can also use gateways on Amazon EC2 instances to provide storage for applications hosted on Amazon EC2. AWS Storage Gateway provides an Amazon Machine Image (AMI) that contains the gateway VM image.

**Topics**

- Step 1: Choose A Gateway Configuration (p. 57)
- Step 2: Deploy and Activate a Gateway on Amazon EC2 (p. 58)

## Step 1: Choose A Gateway Configuration

When you click **Deploy a new Gateway on Amazon EC2**, the setup wizard asks you to first choose a gateway configuration and then launch the AMI.

### To deploy and activate a gateway on an Amazon EC2 instance

1. In the **Setup and Activate Gateway on Amazon EC2** wizard's **Select Gateway Type** section, choose the gateway type you want to activate.

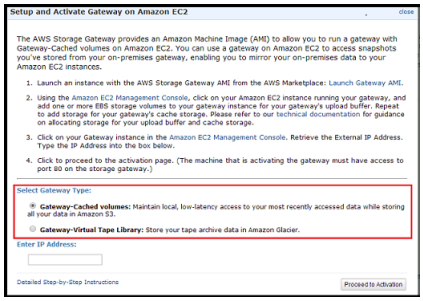
Setup and Activate Gateway on Amazon EC2
close
The AWS Storage Gateway provides an Amazon Machine Image (AMI) to allow you to run a gateway with Gateway-Cached volumes on Amazon EC2. You can use a gateway on Amazon EC2 to access snapshots you've stored from your on-premises gateway, enabling you to mirror your on-premises data to your Amazon EC2 instances.
1. Launch an instance with the AWS Storage Gateway AMI from the AWS Marketplace: Launch Gateway AMI.
2. Using the Amazon EC2 Management Console, click on your Amazon EC2 instance running your gateway, and add one or more EBS storage volumes to your gateway instance for your gateway's upload buffer. Repeat to add storage for your gateway's cache storage. Please refer to our technical documentation for guidance on allocating storage for your upload buffer and cache storage.
3. Click on your Gateway instance in the Amazon EC2 Management Console. Retrieve the External IP Address. Type the IP Address into the box below.
4. Click to proceed to the activation page. (The machine that is activating the gateway must have access to port 80 on the storage gateway.)
Select Gateway Type:
Gateway-Cached volumes: Maintain local, low-latency access to your most recently accessed data while storing all your data in Amazon S3.
Gateway-Virtual Tape Library: Store your tape archive data in Amazon Glacier.
Enter IP Address:
Detailed Step-by-Step Instructions
Proceed to Activation

**Note**
You can only create gateway-cached volumes or a gateway-VTL on an Amazon EC2 instance. For information about gateway types, see How AWS Storage Gateway Works (Architecture) (p. 2).

2. In the **Setup and Activate Gateway on Amazon EC2** wizard, click **Launch Gateway AMI** in step 1. Doing this opens the AWS Marketplace page in a new browser tab.

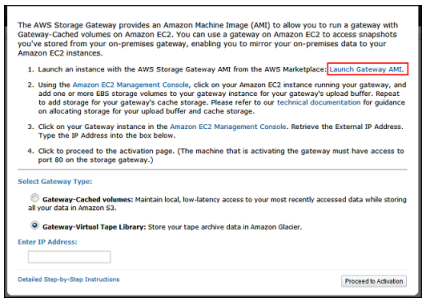
The AWS Storage Gateway provides an Amazon Machine Image (AMI) to allow you to run a gateway with Gateway-Cached volumes on Amazon EC2. You can use a gateway on Amazon EC2 to access snapshots you've stored from your on-premises gateway, enabling you to mirror your on-premises data to your Amazon EC2 instances.
1. Launch an instance with the AWS Storage Gateway AMI from the AWS Marketplace: Launch Gateway AMI.
2. Using the Amazon EC2 Management Console, click on your Amazon EC2 instance running your gateway, and add one or more EBS storage volumes to your gateway instance for your gateway's upload buffer. Repeat to add storage for your gateway's cache storage. Please refer to our technical documentation for guidance on allocating storage for your upload buffer and cache storage.
3. Click on your Gateway instance in the Amazon EC2 Management Console. Retrieve the External IP Address. Type the IP Address into the box below.
4. Click to proceed to the activation page. (The machine that is activating the gateway must have access to port 80 on the storage gateway.)
Select Gateway Type:
Gateway-Cached volumes: Maintain local, low-latency access to your most recently accessed data while storing all your data in Amazon S3.
Gateway-Virtual Tape Library: Store your tape archive data in Amazon Glacier.
Enter IP Address:
Detailed Step-by-Step Instructions
Proceed to Activation

**Next Step**

Step 2: Deploy and Activate a Gateway on Amazon EC2 (p. 58)

# Step 2: Deploy and Activate a Gateway on Amazon EC2

The Amazon Machine Image (AMI) is available on AWS Marketplace (see AWS Storage Gateway) for you to launch an instance either from the Amazon EC2 console or from the AWS Storage Gateway console.

**Important**
Regardless of how you access the AMI, we strongly recommend that you choose the **Launch with EC2 Console** option in AWS Marketplace to launch your instance. You can find the steps for launching this way in the procedure following. If you use the **1-Click Launch** functionality to launch an instance, you need to add Amazon EBS volumes to your instance as a separate step after the instance is launched. For more information, see Adding and Removing Amazon EBS Volumes for Your Gateway-VTL Instance (p. 343).

### Step 1: Choose an AWS Storage Gateway AMI

1. In the **Setup and Activate Gateway on Amazon EC2** wizard, choose your gateway type, and then click **Launch Gateway AMI** to open the AWS Marketplace page for the AMI in your web browser. On that page, click **Continue**.

AWS Storage Gateway
Sold by: Amazon Web Services
The AWS Storage Gateway is a service connecting a software appliance with cloud-based storage to provide seamless and secure integration between an organization's EC2 instances and AWS's storage infrastructure. The service enables you to securely store data to the AWS cloud for scalable and cost-effective storage. The AWS Storage Gateway supports industry-standard storage protocols that work with your existing applications. It provides low-latency performance by maintaining frequently accessed data locally while securely storing all of your data encrypted in the Amazon Simple Storage Service ... Read more
Customer Rating
Be the first to review this product
Latest Version 1.0
Continue

2. On the **Launch on EC2:** page for the AMI, select the **Launch with EC2 Console** tab.

Launch on EC2:
AWS Storage Gateway
1-Click Launch
Review, modify, and launch
Launch with EC2 Console
Info for EC2 Console or API Launches
Click "Accept Terms & Launch with 1-Click" to launch this software with the settings below
Once you accept the terms, you will have access to launch any version of this software in any supported region. For future launches, you can return to this page or launch directly from the EC2 console.

3. If this is your first time using the AWS Storage Gateway AMI, accept the terms of service by clicking **Accept Terms**; otherwise, skip to the next step.

   Keep the browser page open. Within a few moments, a subscription confirmation email is sent to the email address of the account with which you logged in to AWS Marketplace.

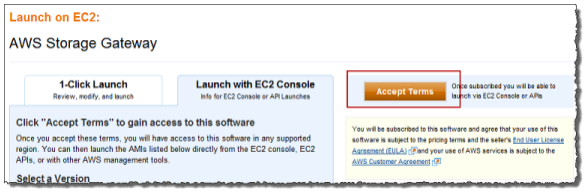
Launch on EC2:
AWS Storage Gateway
1-Click Launch
Review, modify, and launch
Launch with EC2 Console
Info for EC2 Console or API Launches
Accept Terms
Click "Accept Terms" to gain access to this software
Once you accept these terms, you will have access to this software in any supported region. You can then launch the AMIs listed below directly from the EC2 console, EC2 APIs, or with other AWS management tools.
Select a Version

4. In the **Region** list, select the AWS region you want to launch the instance in by clicking the **Launch with EC2 Console** link next to the region name. Note that gateway-VTL AMIs are not available in all regions.

| Region Name | Region String | Gateway-Cached | Gateway-VTL |
|---|---|---|---|
| US East (N. Virginia) | `us-east-1` | yes | yes |

| Region Name | Region String | Gateway-Cached | Gateway-VTL |
|---|---|---|---|
| US West (N. California) | `us-west-1` | yes | yes |
| US West (Oregon) | `us-west-2` | yes | yes |
| EU (Ireland) | `eu-west-1` | yes | yes |
| EU (Frankfurt) | `eu-central-1` | yes | yes |
| Asia Pacific (Singapore) | `ap-southeast-1` | yes | no |
| Asia Pacific (Sydney) | `ap-southeast-2` | yes | yes |
| Asia Pacific (Tokyo) | `ap-northeast-1` | yes | yes |
| South America (Sao Paulo) | `sa-east-1` | yes | no |

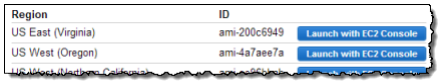


Now the wizard directs you to the Amazon EC2 console for the next step in deploying the AMI.

### Step 2: Choose an Instance Type

1. On the **Choose an Instance Type** page, filter the instances by **All instances** and **All generations** to show the available instance types, and then choose an instance. For this exercise, you can choose the **m1.xlarge** instance type.

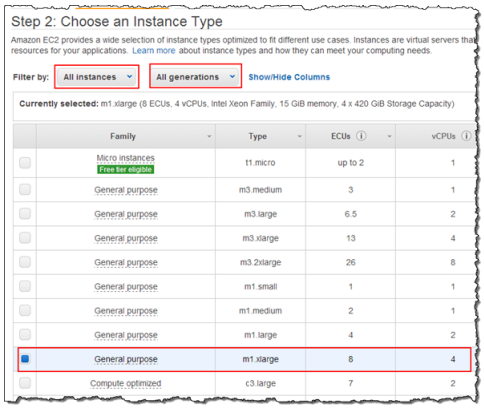


**Note**
AWS Storage Gateway only supports the **m** and **c1** instance types. The instance size must be at least **xlarge**. For more information, go to AWS Storage Gateway in AWS Marketplace.

2. Click **Next: Configure Instance Details** at the bottom of the page.

### Step 3: Configure Instance Details

1. On the **Configure Instance Details** page, you have several configuration options for the instance, including how many instances you want to launch from the AMI and whether you want to launch it

in EC2-Classic or in VPC. If you are creating a gateway for testing purposes, you can keep the default values and go to the next step.

> **Note**
> If your AWS account is a default VPC account, the **Network** drop-down list has a default VPC subnet setting, so that your gateway launches into a default VPC. In that case, you do not have the option to launch into EC2-Classic. If you have an EC2-Classic account, you can create a VPC for your account and launch the instance into the VPC. For information about how to create a VPC, go to Getting Started with Amazon VPC in the *Amazon VPC Getting Started Guide*.

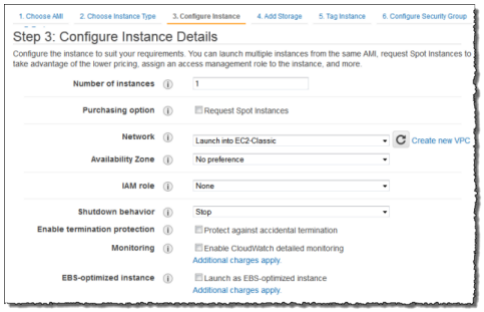


2. Click **Next: Add Storage**.

## Step 4: Add Storage

The Amazon EC2 AMI launches with one root storage volume and a 10GiB EBS volume that is used as a swap disk. You must add additional storage volumes for the gateway to use as cache storage and an upload buffer. For more information, see Volume Gateway Architecture (Gateway-Cached and Gateway-Stored) (p. 2).

> **Important**
> Do not delete any of the existing EBS volumes.

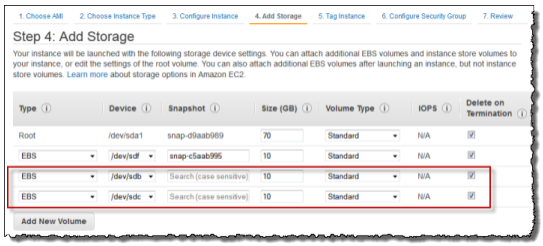


1. Click **Add New Volume** to add additional EBS or instance store volumes.
2. In the **Type** column, select **EBS** or **Instance Store**. The default is EBS.

   > **Note**
   > Instance store volumes are ephemeral—that is, the data on these volumes does not persist when you stop and start or terminate the Amazon EC2 instance. For more information about instance store volumes, go to Amazon EC2 Instance Store in the *Amazon EC2 User Guide*.

3. In the **Size(GiB)** column, enter the size of the volume. If you are creating this gateway as an example setup, you can enter 10 GiB for an upload buffer, and 10 GiB for cache storage. Otherwise, you must size your volumes for a production workload. For more information about sizing these two storage types, see Managing the Upload Buffer and Cache Storage (Gateway-Cached) (p. 169).

The size of the instance store you can add depends on the Amazon EC2 instance type. For more information, see the Instance Types section in Amazon EC2 Instance Store in the *Amazon EC2 User Guide*.

4. Repeat steps 1 through 3 to add another EBS or instance store volume. After this step, you should have a total of four volumes—the root volume, the 10GiB swap disk, and the two volumes you just added.
5. Click **Next: Tag Instance**.

### Step 5: Tag an Instance

1. On the **Tag Instance** page, define a tag with a friendly name for your instance (for example, `Name=MyGatewayVTLOnEC2`). This tag then appears in the Amazon EC2 console where all your Amazon EC2 instances appear.

1. Choose AMI
2. Choose Instance Type
3. Configure Instance
4. Add Storage
5. Tag Instance
6. Configure Security Group
Step 5: Tag Instance
A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. Learn more about tagging your Amazon EC2 resources.
Key (127 characters maximum)
Value (255 characters maximum)
Name
GatewayVTLOnEC2Example
Create Tag
(Up to 10 tags maximum)

2. Click **Next: Configure Security Group**.

### Step 6: Configure a Security Group

The security group must allow port 80 traffic, at least, for activation to occur. To allow connections to iSCSI storage targets of the gateway, you will need also to allow port 3260 traffic. Before launching your instance, you might want to check your existing security groups or create a new security group for your gateway instance. For more information about the security group requirements, see Configuring Security Groups for Your Amazon EC2 Gateway-VTL Instance (p. 344).

We recommend that the gateway's security group only allow traffic on port 3260 to a specific host (the iSCSI initiator) or to a security group to which the initiator belongs.

- On the **Configure Security Group** page, add the following rules. If you are creating a gateway for testing, you can create a new security group, assign a group name, and add these rules.

1. Choose AMI
2. Choose Instance Type
3. Configure Instance
4. Add Storage
5. Tag Instance
6. Configure Security Group
Step 6: Configure Security Group
A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow Internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. Learn more about Amazon EC2 security groups.
Assign a security group:
Create a new security group
Select an existing security group
Security group name:
SecurityGroupForGatewayVTLOnEC2Example
Description:
Created on Friday, October 11, 2013 5:53:09 PM UTC-7
Protocol
Type
Port Range (Code)
Source
SSH
TCP
22
Anywhere
HTTP
TCP
80
Anywhere
HTTPS
TCP
443
Anywhere
RDP
TCP
3389
Anywhere
Add Rule

#### Caution

If you specify **`Anywhere (0.0.0.0/0)`** in **Source**, you enable all IP addresses to access your instance using Secure Shell (SSH). Doing this is acceptable for a short time in a test environment, but it is unsafe for production environments. In production, authorize only a specific IP address or range of addresses to access your instance.

### Step 7: Review and Launch

1. On the **Review and Launch** page, click **Launch**.
2. In the **Select an existing key pair or create a new key pair** window, choose the appropriate option from the drop-down list. If you create a new key pair, give a name to the key pair and download it before you click **Launch Instances**.

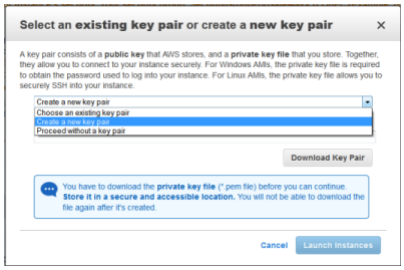


3. In the **Launch Status**, note the instance ID and click **View Instances** to open an Amazon EC2 console.

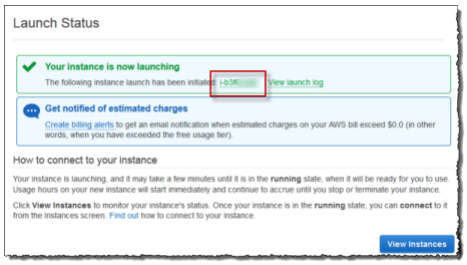
Launch Status
Your instance is now launching
Get notified of estimated charges
How to connect to your instance
View Instances

4. In the Amazon EC2 console, select the Amazon EC2 instance, click the **Description** tab at the bottom, and then note the **IP** address. You will use this IP address to activate the gateway.

   You can use the public or private IP address assigned to the gateway for activation. You must be able to reach the IP address that you use from the browser from which you perform the activation. In this walkthrough, you will use the public IP address to activate the gateway.

### Step 8: Activate the Gateway

1. Now, go back to the AWS Storage Gateway console and provide the IP address in the **Setup and Activate Gateway on Amazon EC2** wizard. Make sure the gateway type you want to activate is selected, and then click **Proceed to Activation**.

   **Note**
   Accessing Amazon EC2 Storage Gateway over a public Internet connection is not supported. The elastic IP address of the Amazon EC2 instance cannot be used as the target address.

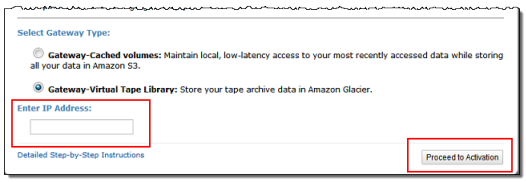
Select Gateway Type:
Gateway-Cached volumes: Maintain local, low-latency access to your most recently accessed data while storing all your data in Amazon S3.
Gateway-Virtual Tape Library: Store your tape archive data in Amazon Glacier.
Enter IP Address:
Detailed Step-by-Step Instructions
Proceed to Activation

2. To complete the activation process, enter the information listed on the activation page. The activation page you see depends on your gateway configuration.

   - **AWS Region** determines where AWS stores your snapshots (for gateway-cached volumes) or virtual tapes (for gateway-VTLs). You select the region in which to deploy your gateway in the **AWS Region** box. You cannot change the region after the gateway is activated.

     **Note**
     AWS Storage Gateway is supported in the following regions.

| Region Name | Region String | Gateway-Cached | Gateway-VTL |
|---|---|---|---|
| US East (N. Virginia) | `us-east-1` | yes | yes |
| US West (N. California) | `us-west-1` | yes | yes |
| US West (Oregon) | `us-west-2` | yes | yes |
| EU (Ireland) | `eu-west-1` | yes | yes |
| EU (Frankfurt) | `eu-central-1` | yes | yes |
| Asia Pacific (Singapore) | `ap-southeast-1` | yes | no |
| Asia Pacific (Sydney) | `ap-southeast-2` | yes | yes |
| Asia Pacific (Tokyo) | `ap-northeast-1` | yes | yes |
| South America (Sao Paulo) | `sa-east-1` | yes | no |

   - **Gateway Time Zone** specifies the time zone to use for your gateway.
   - **Gateway Name** identifies your gateway. You use this name to manage your gateway in the console; you can change it after the gateway is activated. This name must be unique to your account.
   - **Medium Changer Type** specifies the type of medium changer to use for your backup software (available for gateway-VTLs only).

     **Important**
     The type of medium changer you select depends on the backup software you plan to use. The following table shows which medium changer to select for your backup software.

| Backup Software | Medium Changer Type |
|---|---|
| Symantec NetBackup Version 7.x | `AWS-Gateway-VTL` |
| Symantec Backup Exec 2012 | `STK-L700` |
| Symantec Backup Exec 2014 | `AWS-Gateway-VTL` |

| Backup Software | Medium Changer Type |
|---|---|
| Symantec Backup Exec 15 | `AWS-Gateway-VTL` |
| Microsoft System Center 2012 R2 Data Protection Manager | `STK-L700` |
| Veeam Backup & Replication V7 | `STK-L700` |
| Veeam Backup & Replication V8 | `STK-L700` |
| Dell NetVault Backup 10.0 | `STK-L700` |

**Note**
You must select the medium changer that is recommended for your backup software. Other medium changers might not function properly. You can select a different medium changer after the gateway is activated. For more information, see Selecting a Medium Changer After Gateway Activation (p. 424).

- **Tape Drive Type** specifies the type of tape drive used by this gateway (available for gateway-VTLs only).

The following screenshot shows the activation page for gateway-cached volumes or gateway-stored volumes.

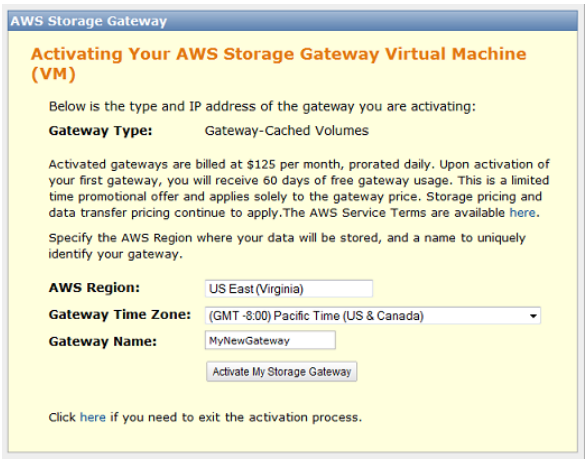
AWS Storage Gateway
Activating Your AWS Storage Gateway Virtual Machine (VM)
Below is the type and IP address of the gateway you are activating:
Gateway Type: Gateway-Cached Volumes
Activated gateways are billed at $125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply.The AWS Service Terms are available here.
Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.
AWS Region: US East (Virginia)
Gateway Time Zone: (GMT -8:00) Pacific Time (US & Canada)
Gateway Name: MyNewGateway
Activate My Storage Gateway
Click here if you need to exit the activation process.

The following screenshot shows the activation page for gateway-VTLs.

AWS Storage Gateway
Activating Your AWS Storage Gateway Virtual Machine
Below is the type of the gateway you are activating:
Gateway Configuration: Gateway-Virtual Tape Library
Activated gateways are billed at $125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply. The AWS Service Terms are available here.
Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.
AWS Region:
US West (Oregon)
Gateway Time Zone:
(GMT) Western Europe Time, London, Lisbon, Casablanca
Gateway Name:
MyGateway-VTL
Medium Changer Type:
AWS-Gateway-VTL
Learn More
Tape Drive Type:
IBM-ULT3580-TD5
Activate My Storage Gateway
Click here if you need to exit the activation process.

3. After you successfully activate the gateway, it appears in your AWS Storage Gateway console in either the **VTL Gateways** or **Volume Gateways** section, depending on the type of gateway you activated.

4. Click the gateway you just deployed. Depending on the type of gateway you activated, the **Create Volumes** or **Create Tapes** button is displayed.

After your Amazon EC2 gateway is activated, you perform the following tasks to set up a fully functional gateway:

- Configure local storage and alarms.
- For a gateway-cached setup, you create volumes. For a gateway-VTL, you create virtual tapes.
- With gateway-cached volumes, you connect the volumes to your Windows client. For a gateway-VTL, you connect your virtual tape library devices.
- Test the setup.

### Next Step

You have completed the deployment and activation steps that are necessary to activate an Amazon EC2 gateway regardless of the gateway type you deploy. The remaining setup steps required to set up functioning gateway are gateway-specific. Click one of the links to continue with Step 2: of the setup process.

- Step 2: Configure Local Storage and Alarms (Gateway-Cached and Gateway-Stored) (p. 67)
- Step 2: Configure Local Storage and Alarms (Gateway-VTL) (p. 241)

# Volume Gateways (Gateway-Cached and Gateway-Stored)

AWS Storage Gateway offers two volume-based storage solutions, gateway-cached volumes and gateway-stored volumes.

The gateway-cached solution lets you create storage volumes and mount them as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers. In this solution, the gateway stores data you write to your gateway-cached volume in Amazon Simple Storage Service (Amazon S3) and stores only a cache of frequently accessed data on your on-premises storage hardware.

The gateway-stored solution lets you store all your data locally in storage volumes on your on-premises storage hardware. In this solution, the gateway periodically takes snapshots as incremental backups and stores them in Amazon S3.

For more information about gateway-cached and gateway-stored volume architecture, see Volume Gateway Architecture (Gateway-Cached and Gateway-Stored) (p. 2). For information about gateway concepts and the storage solutions offered, see What Is AWS Storage Gateway? (p. 1)

The following sections provide information about setting up and managing gateway-cached and gateway-stored volumes. These sections are a continuation of the Getting Started section, and together they provide an end-to-end example setup. You can use the AWS Storage Gateway console to set up and manage volumes, as shown in the following sections. You can also perform these tasks programmatically using the AWS Storage Gateway API or the AWS SDK libraries. For more information on using the AWS Storage Gateway API for this purpose, see API Reference for AWS Storage Gateway (p. 379).

## Related Topics

Getting Started with AWS Storage Gateway (p. 9)

Volume Gateway Architecture (Gateway-Cached and Gateway-Stored) (p. 2)

API Reference for AWS Storage Gateway (p. 379)

# Setting Up Your Volume Gateway (Gateway-Cached and Gateway-Stored)

As you follow the steps in this section, you will use the Setup and Activate Gateway wizard in the AWS Storage Gateway console. At several steps in the wizard, you will perform tasks outside of the console and then return. If your session times out or the browser closes, you can return to the console.

**Topics**


- Step 1: Before You Begin (p. 67)
- Step 2: Configure Local Storage and Alarms (Gateway-Cached and Gateway-Stored) (p. 67)
- Step 3: Create Volumes (p. 74)
- Step 4: Connect Your Volumes to Your Windows Client (p. 79)
- Step 5: Test the Setup (p. 86)
- Step 6: Clean Up After the Example Exercise (p. 94)
- Where Do I Go from Here? (p. 94)


## Step 1: Before You Begin

Before you begin your gateway setup, you need to do the following:

1. Sign up for an AWS account. To use AWS Storage Gateway, you need an AWS account.
2. Deploy and activate a gateway. Steps to deploy and activate a gateway are the same regardless of the type of storage solution you deploy.

You can find instructions for deploying and activating a gateway in the Getting Started section that is common to all gateways, Getting Started with AWS Storage Gateway (p. 9). Before going on, you must first complete these steps. When you have done so, continue to the next step.

**Next Step**

Step 2: Configure Local Storage and Alarms (Gateway-Cached and Gateway-Stored) (p. 67)

## Step 2: Configure Local Storage and Alarms (Gateway-Cached and Gateway-Stored)

Earlier in this Getting Started exercise, you created two disks for your gateway. The disks you created depends on the type of gateway you deployed. Now, you will configure your gateway to use these disks.

- For gateway-cached volumes, you will configure one disk for an upload buffer and the other for cache storage.
- For gateway-stored volumes, you will configure one disk for an upload buffer and allocate the rest of the storage for your application data.

Optionally, you can also create alarms to monitor your disks.

After activation, the **Create Volumes** button appears.

If activation is not successful, see Troubleshooting Your Volume Gateway (Gateway-Cached and Gateway-Stored) (p. 226) for possible solutions.

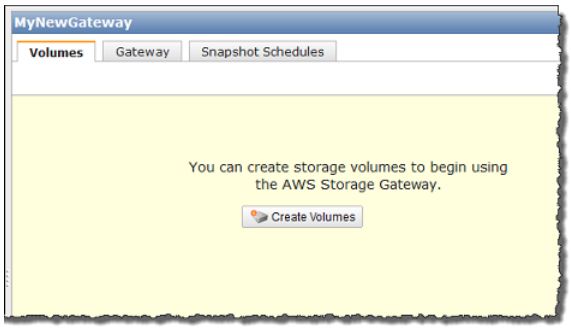


When you click **Create Volumes**, the wizard walks you through how to configure local storage, optionally create alarms, and create volumes for your gateway.

> **Note**
> When creating volumes for a gateway for the first time, you must first configure local storage for the gateway.

**Next Step**

The steps you follow to configure local storage depends on the gateway type you activated. Click one of the following links:

- Configure Local Storage and Alarms (Gateway-Cached Volumes) (p. 68)
- Configure Local Storage and Alarms (Gateway-Stored Volumes) (p. 71)

# Configure Local Storage and Alarms (Gateway-Cached Volumes)

You will now configure one of the two disks for an upload buffer and the other for cache storage, and also optionally create alarms to monitor your disks.

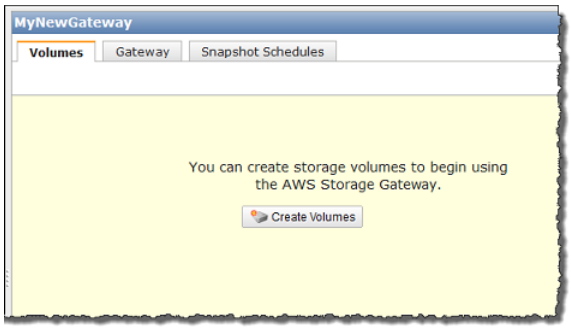


### To configure local storage for an upload buffer and cache storage

1. Click **Create Volumes**.
2. In the **Configure Your Activated Gateway** wizard, configure local working storage.

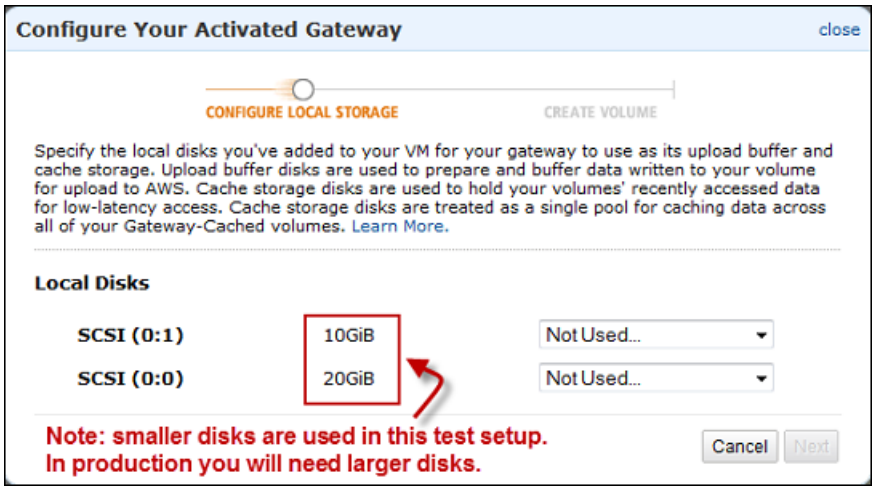
Configure Your Activated Gateway
close
CONFIGURE LOCAL STORAGE
CREATE VOLUME
Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage. Upload buffer disks are used to prepare and buffer data written to your volume for upload to AWS. Cache storage disks are used to hold your volumes' recently accessed data for low-latency access. Cache storage disks are treated as a single pool for caching data across all of your Gateway-Cached volumes. Learn More.
Local Disks
SCSI (0:1)
10GiB
Not Used...
SCSI (0:0)
20GiB
Not Used...
Note: smaller disks are used in this test setup.
In production you will need larger disks.
Cancel
Next

a. Of the two disks you provisioned, select one as the upload buffer and the other as cache storage, as shown in the following screenshot.

Configure Your Activated Gateway
close
CONFIGURE LOCAL STORAGE
CREATE VOLUME
Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage. Upload buffer disks are used to prepare and buffer data written to your volume for upload to AWS. Cache storage disks are used to hold your volumes' recently accessed data for low-latency access. Cache storage disks are treated as a single pool for caching data across all of your Gateway-Cached volumes. Learn More.
Local Disks
SCSI (0:1)
10GiB
Use for Upload Buffer
SCSI (0:0)
20GiB
Use for Cache Storage
Cancel
Next

b. Click **Next**.

Optionally, now you can create alarms to monitor the storage use of the two disks you created.

### To configure optional alarms

1. In the page for the upload buffer alarm, configure the alarms to notify you about upload buffer use.

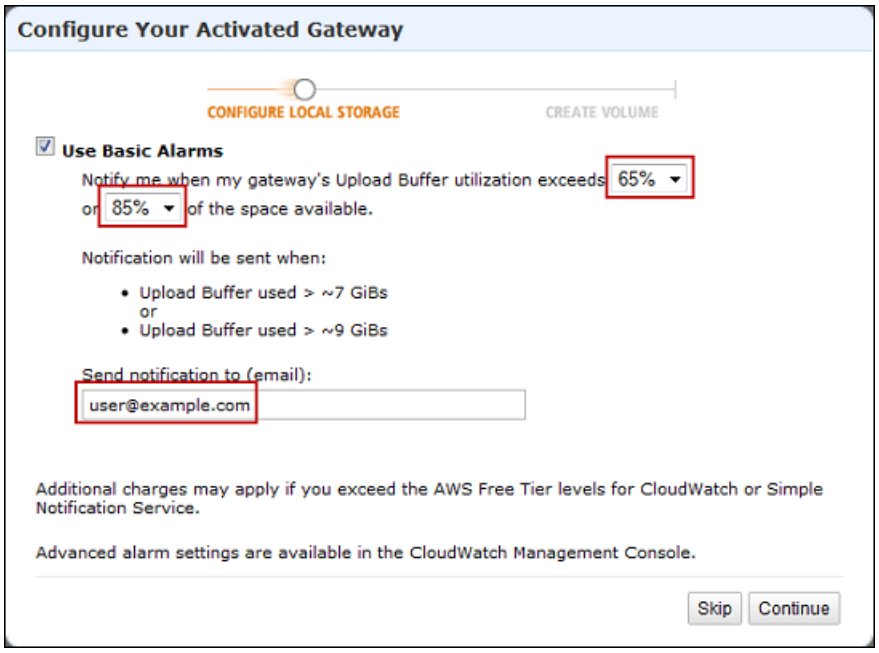
Configure Your Activated Gateway
CONFIGURE LOCAL STORAGE
CREATE VOLUME
Use Basic Alarms
Notify me when my gateway's Upload Buffer utilization exceeds 65%
or 85% of the space available.
Notification will be sent when:
Upload Buffer used > ~7 GiBs
or
Upload Buffer used > ~9 GiBs
Send notification to (email):
user@example.com
Additional charges may apply if you exceed the AWS Free Tier levels for CloudWatch or Simple Notification Service.
Advanced alarm settings are available in the CloudWatch Management Console.
Skip
Continue

a. Using the two drop-down lists, select utilization percentages to set two upload buffer alarms.

   For example, you can set the first threshold (the lower percentage value) to a percentage use of the upload buffer that, if exceeded, you want to be warned about. You can set the second threshold (the higher percentage value) to a percentage use of the upload buffer that, if exceeded, is cause for action. The action might be adding more upload buffer space.

   After you complete this step, you can go to the CloudWatch console at any time and change the alarm thresholds.

b. Type an email address to which alarm notifications should be sent.

c. Click **Continue**.

   Two alarms are created. For example, if you use the gateway name *MyNewGatewayCached*, the alarms created are *MyNewGatewayCached-UploadBufferUtilization-Alarm1* and *MyNewGatewayCached-UploadBufferUtilization-Alarm2*.

d. Check the email address that you provided for a subscription confirmation email, and follow the instructions in that email to confirm your subscription to the related Amazon Simple Notification Service (Amazon SNS) topic. After you have confirmed your subscription, you will receive an email when either threshold you specified is exceeded.

2. Click **Continue** to create the alarms.
3. In the cached storage alarms page, configure the alarms to notify you about cache storage use.

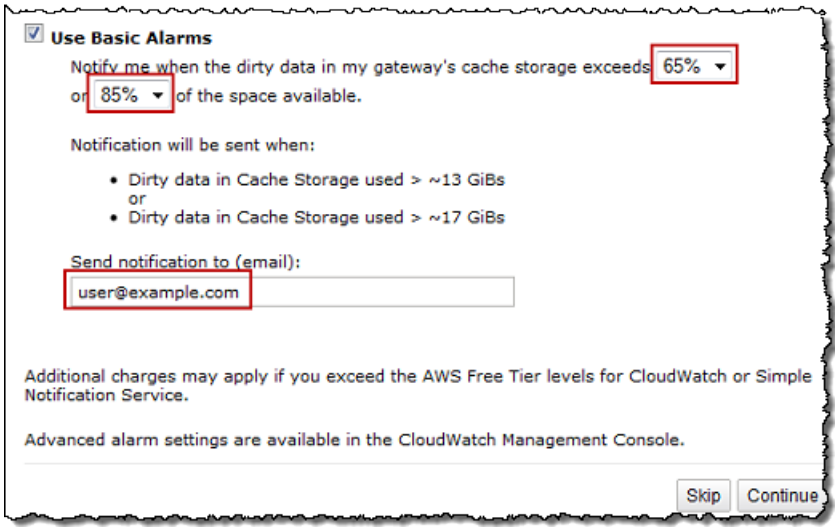
Use Basic Alarms
Notify me when the dirty data in my gateway's cache storage exceeds 65%
or 85% of the space available.
Notification will be sent when:
Dirty data in Cache Storage used > ~13 GiBs
or
Dirty data in Cache Storage used > ~17 GiBs
Send notification to (email):
user@example.com
Additional charges may apply if you exceed the AWS Free Tier levels for CloudWatch or Simple Notification Service.
Advanced alarm settings are available in the CloudWatch Management Console.
Skip
Continue

Follow the instructions in the preceding step to configure these alarms.

For example, if you use the gateway name *MyNewGatewayCached*, the alarms created are *MyNewGatewayCached-CacheUtilization-Alarm1* and *MyNewGatewayCached-CacheUtilization-Alarm2*.

4. Click **Continue** to create the alarms.

**Next Step**

Create Volumes for a Gateway-Cached Configuration (p. 75)

# Configure Local Storage and Alarms (Gateway-Stored Volumes)

You will now configure one of the two disks for use by your application's data and the other disk as an upload buffer, and also optionally create alarms to monitor your disks.

> **Note**
> When you configure a disk as an upload buffer, you lose any existing data on the disk, so be careful to preserve your data.

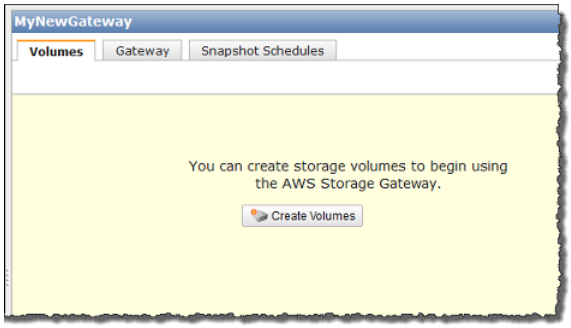


### To configure local storage for application data and an upload buffer

1. Click **Create Volumes**.

The **Create Storage Volume** page appears.

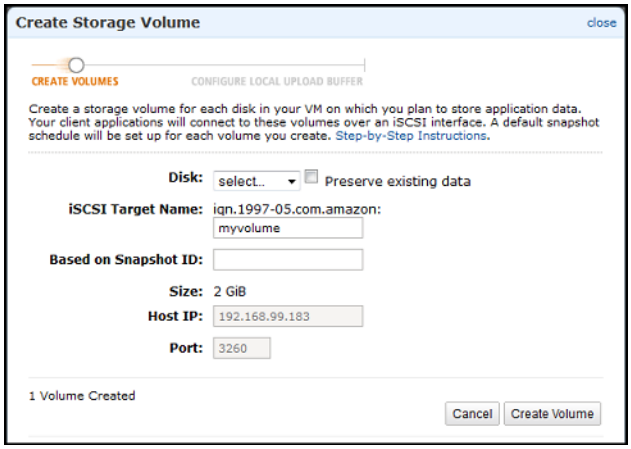
Create Storage Volume
close
CREATE VOLUMES
CONFIGURE LOCAL UPLOAD BUFFER
Create a storage volume for each disk in your VM on which you plan to store application data. Your client applications will connect to these volumes over an iSCSI interface. A default snapshot schedule will be set up for each volume you create. Step-by-Step Instructions.
Disk:
select..
Preserve existing data
iSCSI Target Name:
iqn.1997-05.com.amazon:
myvolume
Based on Snapshot ID:
Size:
2 GiB
Host IP:
192.168.99.183
Port:
3260
1 Volume Created
Cancel
Create Volume

2. In the **Disk** drop-down list, select the disk you want to use to store application data.
3. In the **iSCSI Target Name** text box, type a name for your volume, and then click **Create Volume**.
4. Click **Configure Local Upload Buffer**, at lower right.

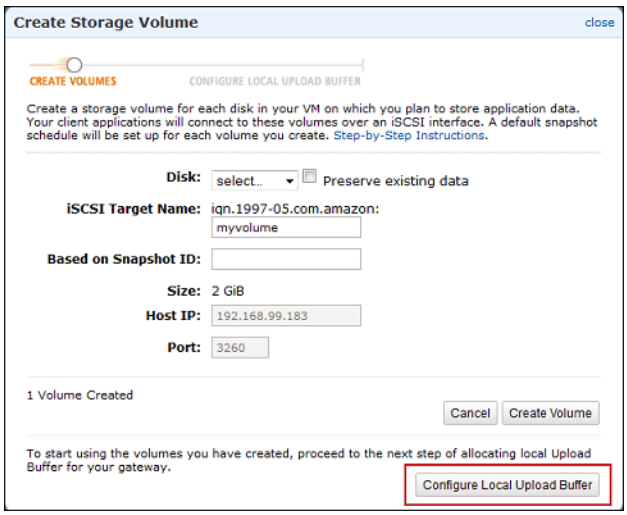
Create Storage Volume
close
CREATE VOLUMES
CONFIGURE LOCAL UPLOAD BUFFER
Create a storage volume for each disk in your VM on which you plan to store application data. Your client applications will connect to these volumes over an iSCSI interface. A default snapshot schedule will be set up for each volume you create. Step-by-Step Instructions.
Disk:
select..
Preserve existing data
iSCSI Target Name:
iqn.1997-05.com.amazon:
myvolume
Based on Snapshot ID:
Size:
2 GiB
Host IP:
192.168.99.183
Port:
3260
1 Volume Created
Cancel
Create Volume
To start using the volumes you have created, proceed to the next step of allocating local Upload Buffer for your gateway.
Configure Local Upload Buffer

5. Click the check box next to the remaining disk to allocate the disk as the upload buffer, and then click **Next**.

   This page lists all available disks on your virtual machine (VM). Earlier, you added two virtual disks to the VM and configured one of the disks as a storage volume. Therefore, the page should show one available disk. Select this disk to allocate as the upload buffer. You can extend the upload buffer later without disrupting iSCSI input/output.

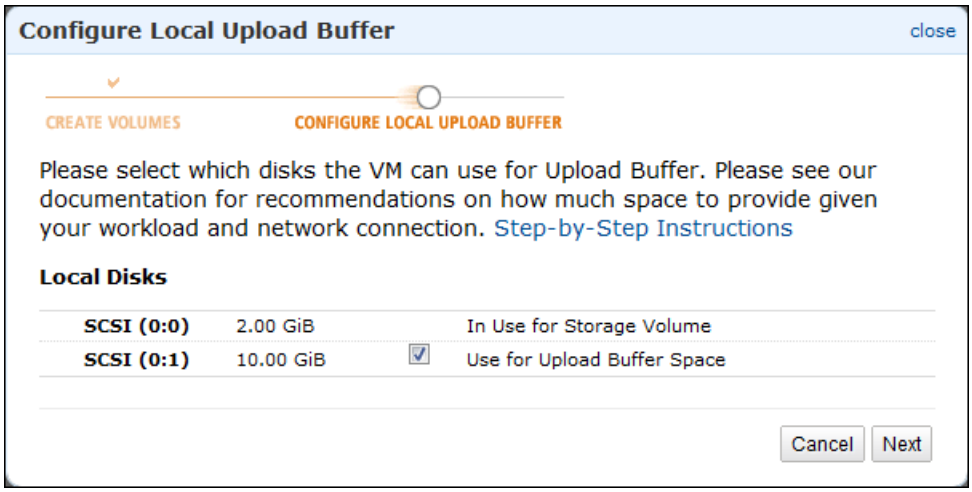
Configure Local Upload Buffer
close
CREATE VOLUMES
CONFIGURE LOCAL UPLOAD BUFFER
Please select which disks the VM can use for Upload Buffer. Please see our documentation for recommendations on how much space to provide given your workload and network connection. Step-by-Step Instructions
Local Disks
SCSI (0:0) 2.00 GiB In Use for Storage Volume
SCSI (0:1) 10.00 GiB Use for Upload Buffer Space
Cancel
Next

6. In the confirmation page, select the check box and click **Confirm**.

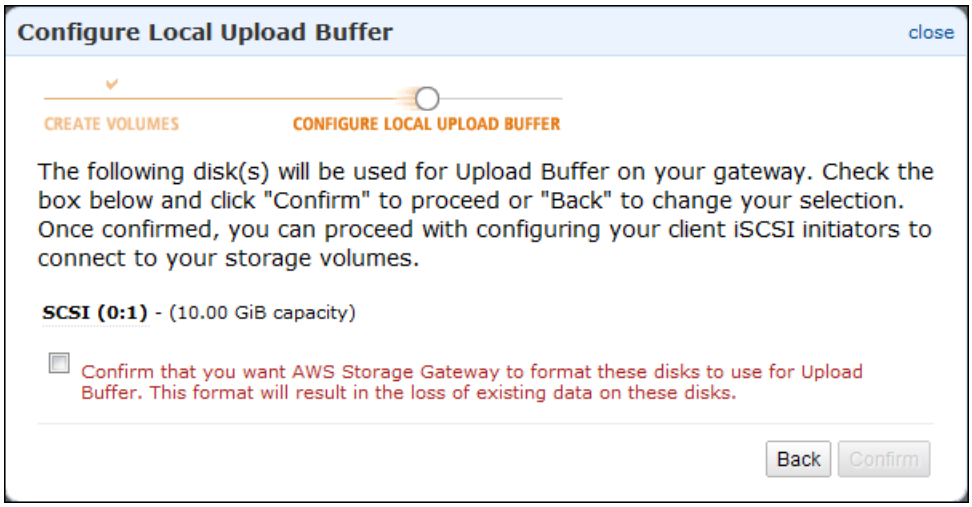
Configure Local Upload Buffer
close
CREATE VOLUMES
CONFIGURE LOCAL UPLOAD BUFFER
The following disk(s) will be used for Upload Buffer on your gateway. Check the box below and click "Confirm" to proceed or "Back" to change your selection. Once confirmed, you can proceed with configuring your client iSCSI initiators to connect to your storage volumes.
SCSI (0:1) - (10.00 GiB capacity)
Confirm that you want AWS Storage Gateway to format these disks to use for Upload Buffer. This format will result in the loss of existing data on these disks.
Back
Confirm

### To configure optional alarms

1. In the **Upload Buffer Alarms** page, configure alarms for your upload buffer.

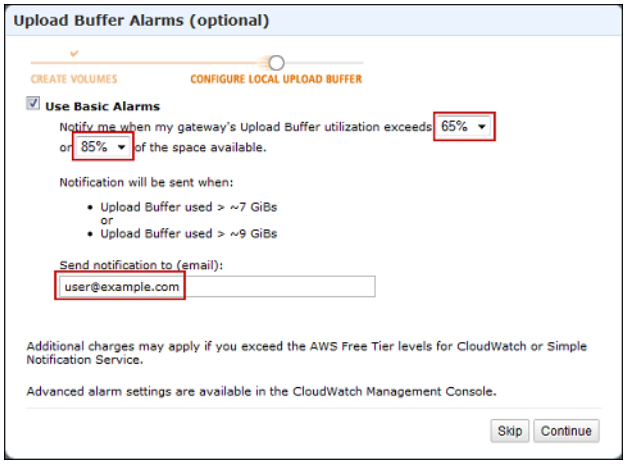
Upload Buffer Alarms (optional)
CREATE VOLUMES
CONFIGURE LOCAL UPLOAD BUFFER
Use Basic Alarms
Notify me when my gateway's Upload Buffer utilization exceeds 65%
or 85% of the space available.
Notification will be sent when:
Upload Buffer used > ~7 GiBs
or
Upload Buffer used > ~9 GiBs
Send notification to (email):
user@example.com
Additional charges may apply if you exceed the AWS Free Tier levels for CloudWatch or Simple Notification Service.
Advanced alarm settings are available in the CloudWatch Management Console.
Skip
Continue

a. Using the two drop-down lists, select utilization percentages to set two upload buffer alarms.

   For example, you can set the first threshold (the lower percentage value) to a percentage use of the upload buffer that, if exceeded, you want to be warned about. You can set the second threshold to a percentage use of the upload buffer that, if exceeded, is cause for action. The action might be adding more upload buffer space.

   After you complete this step, you can go to the CloudWatch console at any time and change the alarm thresholds.

b. Type an email address to which alarm notifications should be sent.

c. Click **Continue**.

   Two alarms are created. For example, if you use the gateway name *MyNewGateway*, the alarms created are *MyNewGateway-UploadBufferUtilization-Alarm1* and *MyNewGateway-UploadBufferUtilization-Alarm2*.

d. Check the email address that you provided for a subscription confirmation email, and follow the instructions in that email to confirm your subscription to the Amazon Simple Notification Service (Amazon SNS). After you have confirmed your subscription, you will receive an email when either threshold you specified is exceeded.

For more detailed information about creating upload buffer alarms, see Monitoring the Upload Buffer (p. 191).

2. In the **Configure iSCSI Initiators** page, click **Close**.

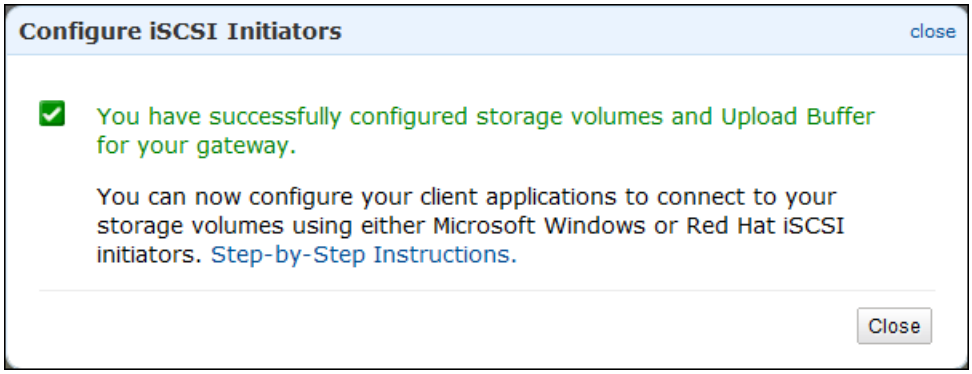
Configure iSCSI Initiators
close
You have successfully configured storage volumes and Upload Buffer for your gateway.
You can now configure your client applications to connect to your storage volumes using either Microsoft Windows or Red Hat iSCSI initiators. Step-by-Step Instructions.
Close

**Next Step**

Create Volumes for a Gateway-Stored Configuration (p. 77)

# Step 3: Create Volumes

Now that you have activated your gateway and configured local storage, you will create storage volumes to store your data. The storage you create depends on the type of gateway you activated.

**Next Step**

Depending on the gateway configuration you activated, click one of the following links:

- Create Volumes for a Gateway-Cached Configuration (p. 75)
- Create Volumes for a Gateway-Stored Configuration (p. 77)

# Create Volumes for a Gateway-Cached Configuration

Previously, you allocated one of two disks that you had added to the VM for cache storage and the other disk for an upload buffer. You will now create a storage volume in Amazon S3 and configure CHAP authentication for your volumes. Your on-premises applications writes data to this volume and stores the data in Amazon S3. The gateway maintains the volume's recently accessed data locally in cache storage.

**Topics**


- Create a Volume in Amazon S3 (p. 75)
- Configure CHAP Authentication for Your Volumes (p. 77)


The following architectural overview diagram shows the part of the gateway-cached setup that you are creating.

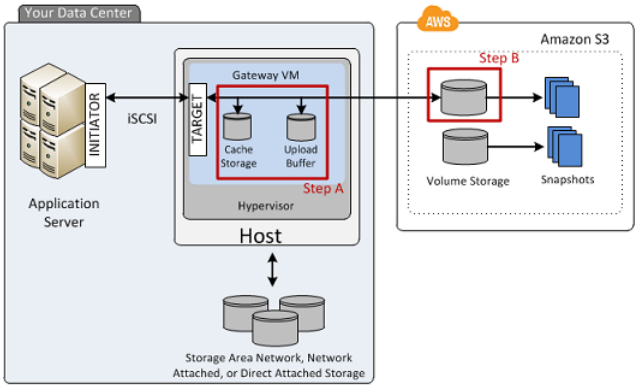


## Create a Volume in Amazon S3

### To create a volume in Amazon S3 for a gateway-cached configuration

1. In the **Configure Your Activated Gateway** dialog box, create an iSCSI storage volume in Amazon S3.

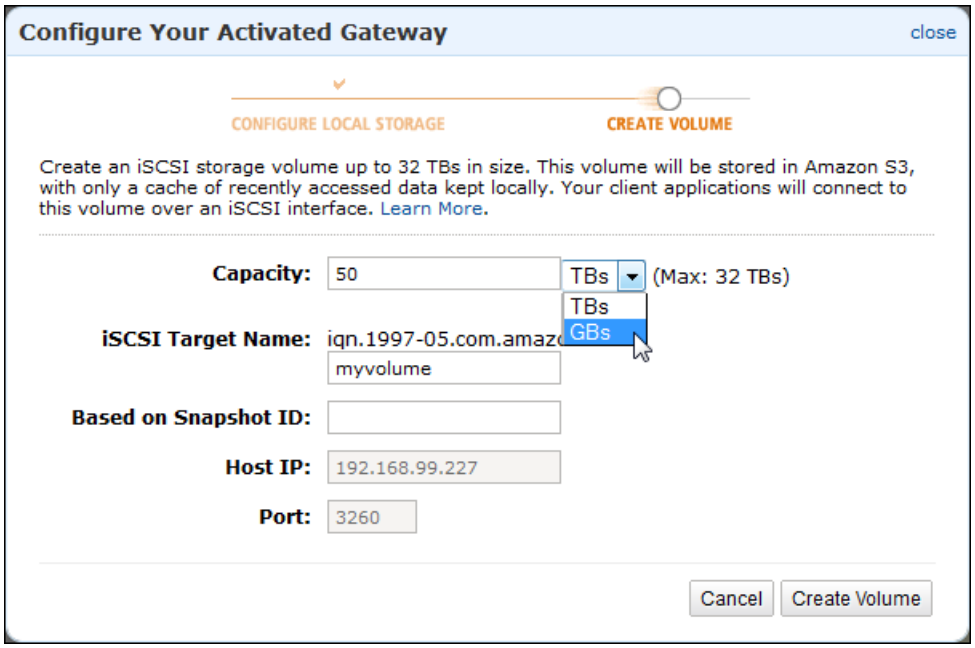


a. For **Capacity**, specify the appropriate capacity to hold the data you plan to store in AWS. If you are creating the gateway as an example setup, you can specify 50 GB.

   The maximum size you can specify is 32 TB.

b. Type a name in the **iSCSI Target Name** box.

   The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** UI after discovery. For example, the name `target1` appears as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your storage area network (SAN).

   For this exercise, you can use **myvolume** as the target name.

c. Leave the **Based on Snapshot ID** box empty.

   If you want to restore an existing Amazon EBS snapshot or a gateway snapshot on the storage volume that you are creating, you must specify the snapshot ID. The gateway downloads your existing snapshot data to the storage volume. For more information, see Restoring a Snapshot to a Storage Volume (p. 222).

d. Verify that the **Host IP** setting is the IP address of your gateway, and then click **Create Cached Volume**.

2. In the **Configure iSCSI Initiators** dialog box, click **Close**.

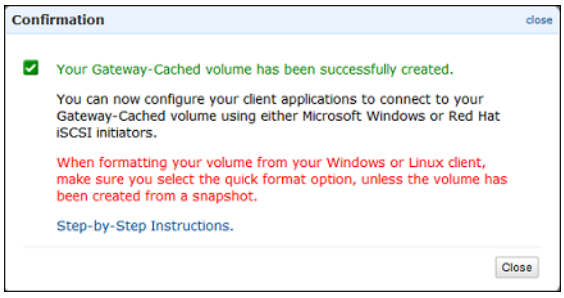
Confirmation
close
Your Gateway-Cached volume has been successfully created.
You can now configure your client applications to connect to your Gateway-Cached volume using either Microsoft Windows or Red Hat iSCSI initiators.
When formatting your volume from your Windows or Linux client, make sure you select the quick format option, unless the volume has been created from a snapshot.
Step-by-Step Instructions.
Close

For more information, see Adding a Storage Volume (p. 97).

The Configure CHAP Authentication wizard appears.

## Configure CHAP Authentication for Your Volumes

You can configure your volumes to use Challenge-Handshake Authentication Protocol (CHAP). CHAP provides protection against playback attacks by requiring authentication to access your storage volume targets.

In the Configure CHAP Authentication dialog box, you provide information so that you can configure CHAP for your volumes.

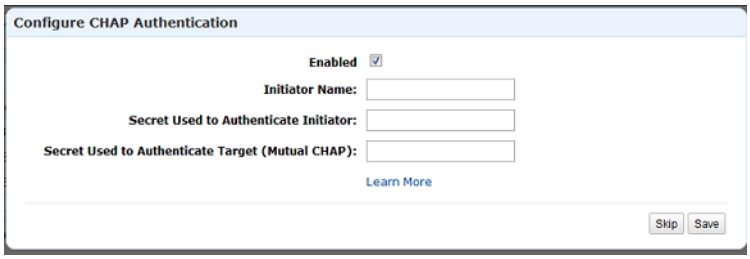


### To configure CHAP

1. In the Configure CHAP Authentication dialog box, click the **Enabled** check box.
2. In the **Initiator Name** box, type the name of your initiator.
3. In the **Secret Used to Authenticate Initiator** box, type the secret phrase you used to authenticate your iSCSI initiator.
4. In the **Secret Used to Authenticate Target (Mutual CHAP)** box, type the secret phrase used to authenticate your target for mutual CHAP.
5. Click **Save** to save these values.

   For more details on setting up CHAP authentication, see Configuring CHAP Authentication for Your Storage Volume (p. 116).

### Next Step

Step 4: Connect Your Volumes to Your Windows Client (p. 79)

# Create Volumes for a Gateway-Stored Configuration

Previously, you allocated one of the two disks that you added to the VM for your application's data and the other disk for an upload buffer. You will now create a storage volume for your application data and configure CHAP authentication for your volumes. Your on-premises applications write data to this storage volume and store the data locally. The gateway periodically takes snapshots (incremental backups) and uploads them to Amazon S3.

### Topics

- Create a Storage Volume (p. 78)
- Configure CHAP Authentication for Your Volumes (p. 79)

The following architectural overview diagram shows the part of the gateway-stored setup that you are creating.

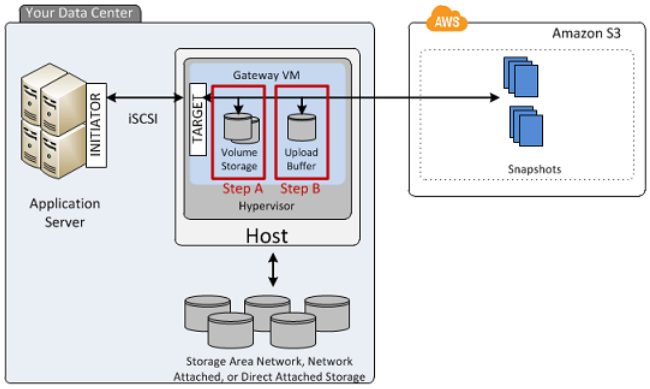


## Create a Storage Volume

Your application data will reside in the volume that you create.

### To create a storage volume for a gateway-stored configuration

- In the **Create Storage Volume** wizard, provide the storage volume information.

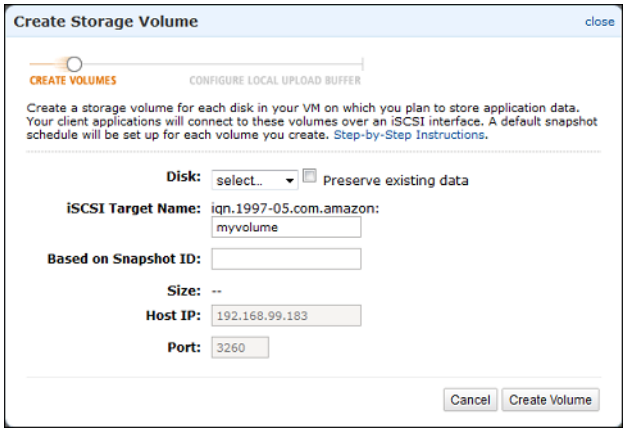
Create Storage Volume
close
CREATE VOLUMES
CONFIGURE LOCAL UPLOAD BUFFER
Create a storage volume for each disk in your VM on which you plan to store application data. Your client applications will connect to these volumes over an iSCSI interface. A default snapshot schedule will be set up for each volume you create. Step-by-Step Instructions.
Disk:
select..
Preserve existing data
iSCSI Target Name:
iqn.1997-05.com.amazon:
myvolume
Based on Snapshot ID:
Size:
--
Host IP:
192.168.99.183
Port:
3260
Cancel
Create Volume

  a. In the **Disk** drop-down list, select the 150 GiB virtual disk on your VM.

     This drop-down list shows the virtual disks that you added to the gateway VM. Select the disk on which you plan to store data.

  b. Keep the **Preserve existing data** check box unchecked.

     **Caution**
     Make sure that you don't have any existing data on the virtual disk. Any existing data on the disk is lost.

  c. Type a name in the **iSCSI Target Name** box.

     The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** UI after discovery. For example, the name `target1` appears as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN.

  d. Leave the **Based on Snapshot ID** box empty.

If you want to restore an existing Amazon EBS snapshot or a gateway snapshot on the storage volume that you are creating, you must specify the snapshot ID. The gateway downloads your existing snapshot data to the storage volume. For more information, see Restoring a Snapshot to a Storage Volume (p. 222).

e. Verify that the **Host IP** setting is the IP address of your gateway, and then click **Create Volume**.

## Configure CHAP Authentication for Your Volumes

You can configure your volumes to use Challenge-Handshake Authentication Protocol (CHAP). CHAP provides protection against playback attacks by requiring authentication to access your storage volume targets.

In the Configure CHAP Authentication dialog box, you provide information so that you can configure CHAP for your volumes.

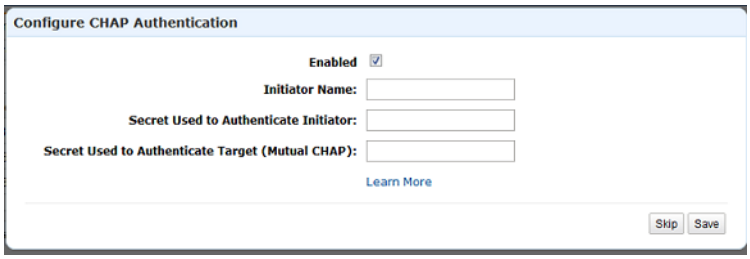


### To configure CHAP

1. In the Configure CHAP Authentication dialog box, click the **Enabled** check box.
2. In the **Initiator Name** box, type the name of your initiator.
3. In the **Secret Used to Authenticate Initiator** box, type the secret phrase you used to authenticate your iSCSI initiator.
4. In the **Secret Used to Authenticate Target (Mutual CHAP)** box, type the secret phrase used to authenticate your target for mutual CHAP.
5. Click **Save** to save these values.

   For more details on setting up CHAP authentication, see Configuring CHAP Authentication for Your Storage Volume (p. 116).

### Next Step

Step 4: Connect Your Volumes to Your Windows Client (p. 79)

# Step 4: Connect Your Volumes to Your Windows Client

In this step, you use the Windows iSCSI initiator to connect to your gateway's volumes. At the end of this procedure, the volumes become available as local devices on your Windows client. For instructions on accessing the iSCSI storage volume from Linux, see Connecting from a Red Hat Client to Your Storage Volume (p. 114).

AWS Storage Gateway enables applications that are clustered using Windows Server Failover Clustering (WSFC) to use the iSCSI initiator to access your gateway's volumes. However, connecting multiple hosts to the same iSCSI target is not supported.

**Note**
You can enable Challenge-Handshake Authentication Protocol (CHAP) authentication for your iSCSI storage volumes. For information about how to enable CHAP, see Configuring CHAP Authentication for Your Storage Volume (p. 116).

### To connect your Windows client to a storage volume

1. On the **Start** menu of your Windows client computer, type **`iscsicpl.exe`** in the **Search Programs and files** box, locate the iSCSI initiator program and then run it.

   **Note**
   You must have administrator rights on the client computer to run the iSCSI initiator.

2. If prompted, click **Yes** to start the Microsoft iSCSI initiator service.

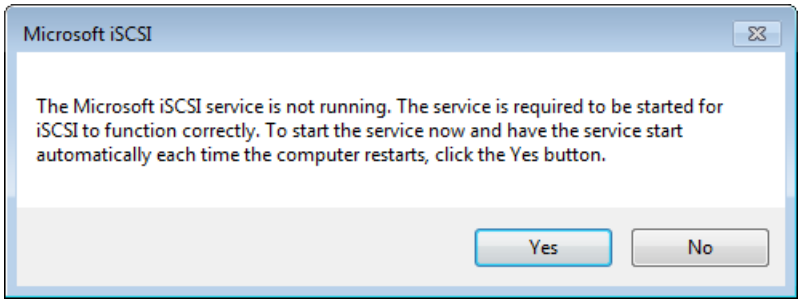
Microsoft iSCSI
The Microsoft iSCSI service is not running. The service is required to be started for iSCSI to function correctly. To start the service now and have the service start automatically each time the computer restarts, click the Yes button.
Yes
No

3. In the **iSCSI Initiator Properties** dialog box, click the **Discovery** tab, and then click the **Discovery Portal** button.

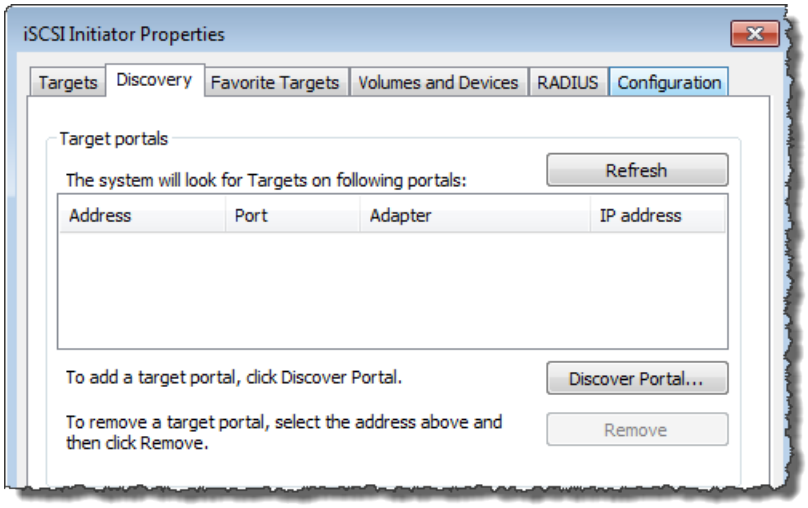


4. In the **IP address or DNS name** text box of the **Discover Target Portal** dialog box, type the IP address of your iSCSI target, and then click **OK**. You can obtain the IP address of your gateway in the **Gateway** tab on the AWS Storage Gateway console. If you deployed your gateway on an Amazon EC2 instance, you can obtain the public IP address or public DNS in the **Description** tab on the Amazon EC2 console.

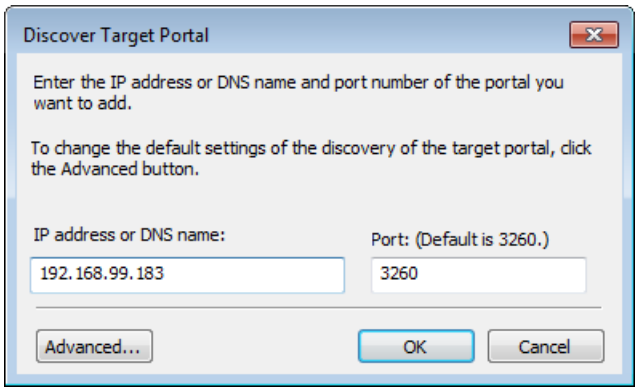
Discover Target Portal
Enter the IP address or DNS name and port number of the portal you want to add.
To change the default settings of the discovery of the target portal, click the Advanced button.
IP address or DNS name:
Port: (Default is 3260.)
192.168.99.183
3260
Advanced...
OK
Cancel

The IP address now appears in the list of **Target portals** in the **Discovery** tab.

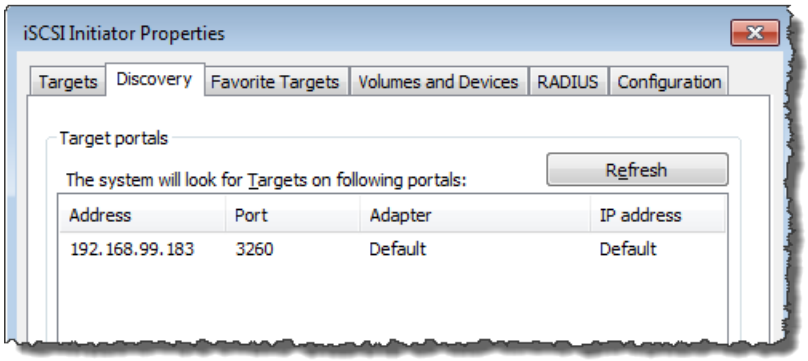


5. Connect the new target portal to the storage volume target on the gateway:

   a. Click the **Targets** tab.

      The new target portal is shown with an inactive status. Note that the target name shown should be the same as the name you specified for your storage volume in step 1.

   b. Select the target, and then click **Connect**.

      If the target name is not populated already, type the name of the target as shown in step 1 in the **Connect to Target** dialog box, select the check box next to **Add this connection to the list of Favorite Targets**, and then click **OK**.

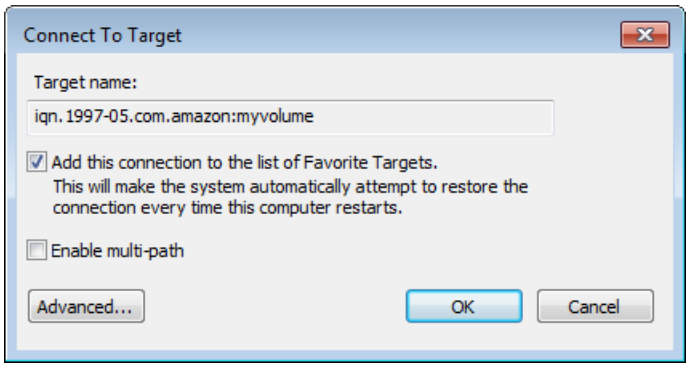
Connect To Target
Target name:
iqn.1997-05.com.amazon:myvolume
Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.
Enable multi-path
Advanced...
OK
Cancel

c. In the **Targets** tab, ensure that the target **Status** has the value **Connected** indicating the target is connected, and then click **OK**.

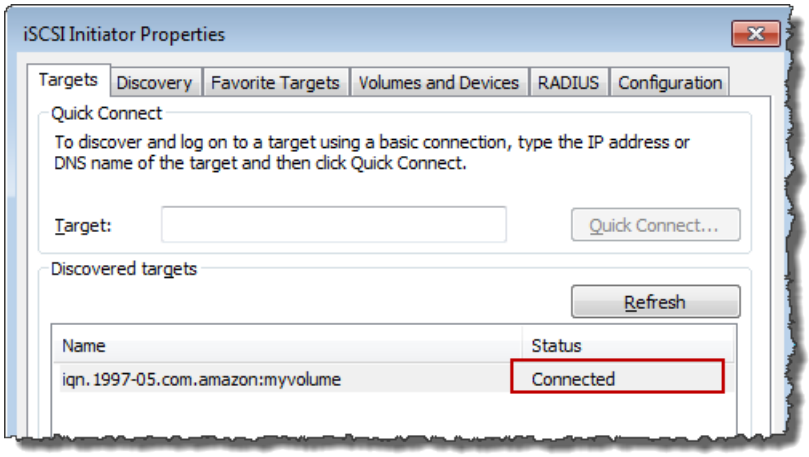


You can now initialize and format this storage volume for Windows so you can begin saving data on it. You do this by using the Windows Disk Management tool.

**Note**
Although it is not required for this exercise, we highly recommend that you customize your iSCSI settings for a real-world application as discussed in the topic Customizing Your Windows iSCSI Settings (p. 112).

### To initialize and format the storage volume you just mapped

1. In the **Start** menu, type **`diskmgmt.msc`** to open the **Disk Management** console.
2. In the **Initialize Disk** dialog box, select **MBR (Master Boot Record)** as the partition style, and then click **OK**. When selecting the partition style, you should take into account the type of volume you are connecting to—gateway-cached or gateway-stored—as shown in the following table.

| Partition Style | Use in the Following Conditions |
|---|---|
| **MBR (Master Boot Record)** | • If your gateway is a gateway-stored volume and the storage volume is limited to 1 TiB in size.<br>• If your gateway is a gateway-cached volume and the storage volume is less than 2 TiB in size. |

| Partition Style | Use in the Following Conditions |
|---|---|
| **GPT (GUID Partition Table)** | If your gateway is a gateway-stored volumes or gateway-cached volumes and the storage volume is 2 TiB or greater in size. |

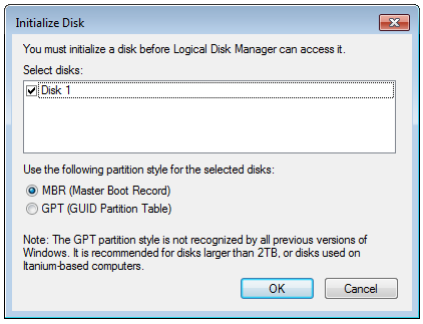
Initialize Disk
You must initialize a disk before Logical Disk Manager can access it.
Select disks:
Disk 1
Use the following partition style for the selected disks:
MBR (Master Boot Record)
GPT (GUID Partition Table)
Note: The GPT partition style is not recognized by all previous versions of Windows. It is recommended for disks larger than 2TB, or disks used on Itanium-based computers.
OK
Cancel

3. Create a simple volume:

   a. If the disk is offline, you must bring it online before you can initialize it. After the disk is initialized, you can format it as a simple volume. All the available volumes are displayed in the disk management console. In the following example, **Disk 1** is the storage volume. Notice that when you select the new volume, it displays hatch lines indicating that it is selected.

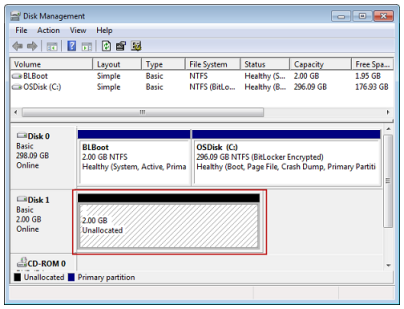


   b. Right-click the disk and select **New Simple Volume**.

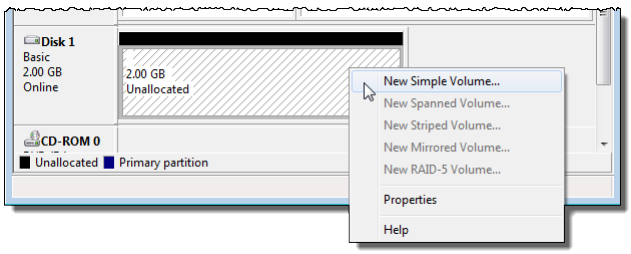


**Important**
Be careful not to format the wrong disk. Check to ensure that the disk you are formatting matches the size of the local disk you allocated to the gateway VM and that it has a status of **Unallocated**.

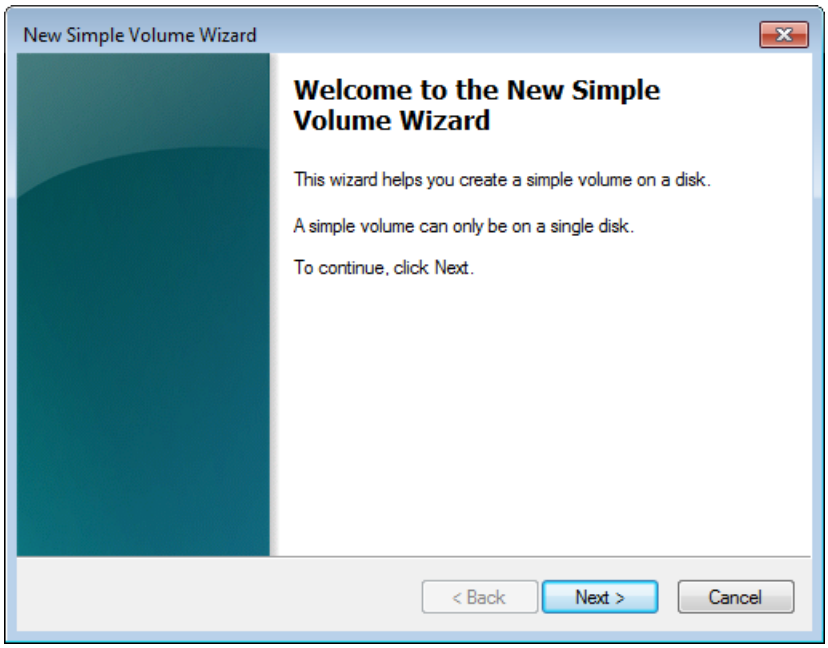
New Simple Volume Wizard
Welcome to the New Simple Volume Wizard
This wizard helps you create a simple volume on a disk.
A simple volume can only be on a single disk.
To continue, click Next.
< Back
Next >
Cancel

c. In the **New Simple Volume Wizard**, click **Next**.

d. In the **Specify Volume Size** dialog box, keep the default values, and then click **Next**.

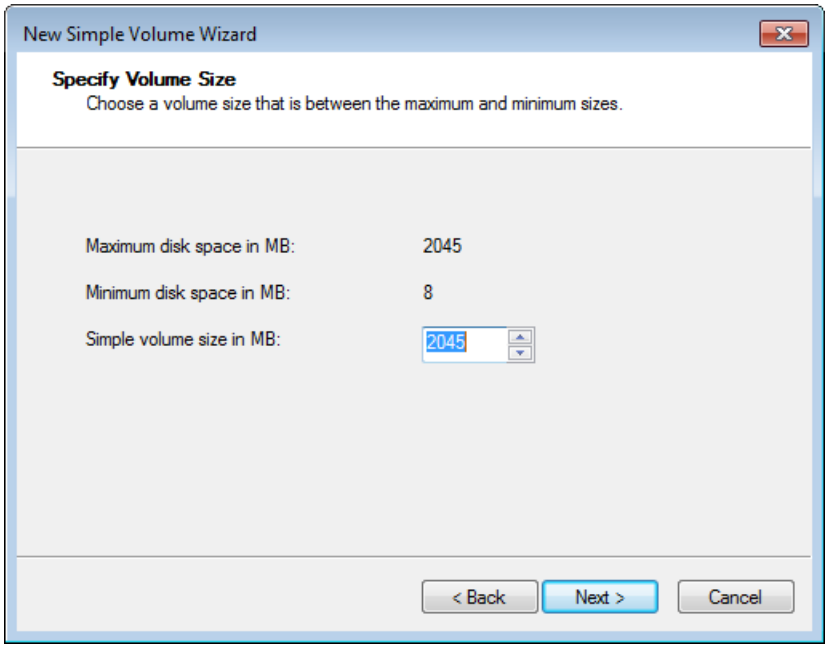


e. In the **Assign Drive Letter or Path** dialog box, keep the default values, and then click **Next**.

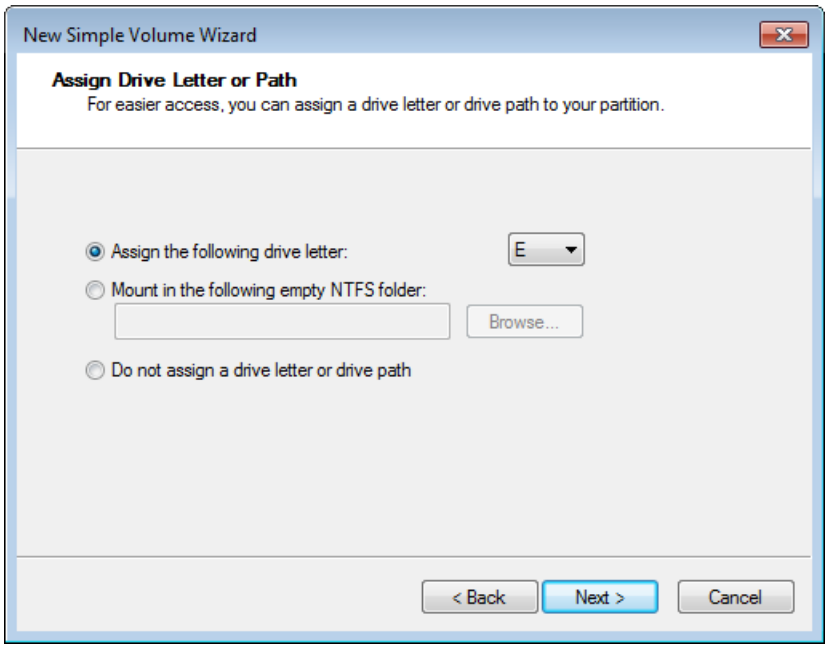
New Simple Volume Wizard
Assign Drive Letter or Path
For easier access, you can assign a drive letter or drive path to your partition.
Assign the following drive letter:
E
Mount in the following empty NTFS folder:
Browse...
Do not assign a drive letter or drive path
< Back
Next >
Cancel

f. In the **Format Partition** dialog box, type a label in the **Volume label** box, and ensure that **Perform a quick format** is selected. Click **Next**.

> **Caution**
> Selecting **Perform a quick format** is highly recommended for gateway-cached volumes because it results in less initialization I/O, smaller initial snapshot size, and the fastest time to a usable volume. It also avoids gateway-cached volume usage that comes from the full format process and not from application data activity.

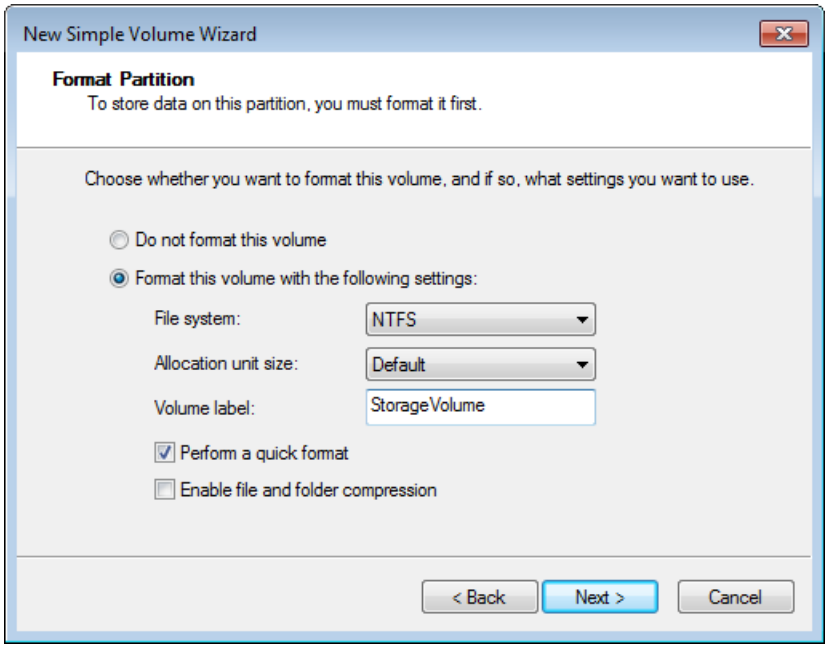
New Simple Volume Wizard
Format Partition
To store data on this partition, you must format it first.
Choose whether you want to format this volume, and if so, what settings you want to use.
Do not format this volume
Format this volume with the following settings:
File system:
NTFS
Allocation unit size:
Default
Volume label:
StorageVolume
Perform a quick format
Enable file and folder compression
< Back
Next >
Cancel

g. Click **Finish** to close the wizard.

> **Note**
> The time that it takes to format the volume depends on the size of the volume. The process might take several minutes to complete.

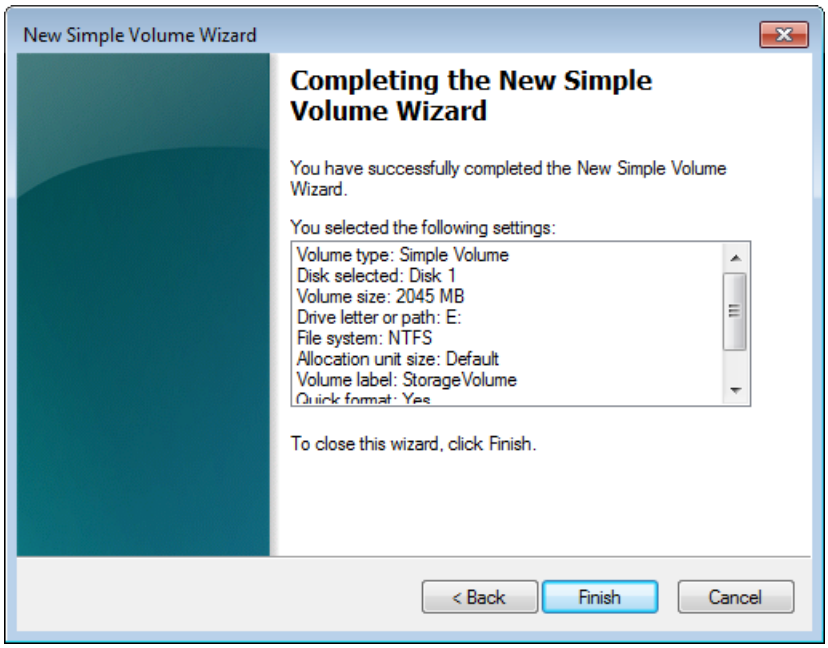
New Simple Volume Wizard
Completing the New Simple Volume Wizard
You have successfully completed the New Simple Volume Wizard.
You selected the following settings:
Volume type: Simple Volume
Disk selected: Disk 1
Volume size: 2045 MB
Drive letter or path: E:
File system: NTFS
Allocation unit size: Default
Volume label: StorageVolume
Quick format: Yes
To close this wizard, click Finish.
< Back
Finish
Cancel

**Next Step**

Step 5: Test the Setup (p. 86)

# Step 5: Test the Setup

You test your volume gateway setup by performing the following tasks:

1. Write data to the volume.
2. Take a snapshot.
3. Restore the snapshot to another volume.

**Next Step**

Depending on the type of gateway you activated, click one of the following links.

- Test the Setup for a Gateway-Cached Volume (p. 86)
- Test the Setup for a Gateway-Stored Volume (p. 90)

## Test the Setup for a Gateway-Cached Volume

You verify the setup for a gateway-cached volume by taking a snapshot backup of the volume and storing the snapshot in AWS. You then restore the snapshot to new volume. Your gateway copies the data from the specified snapshot in AWS to the new volume.

> **Note**
> Restoring data from Amazon EBS volumes that are encrypted is not supported.

### To create a snapshot of a storage volume

1. On your Windows computer, copy some data to your mapped storage volume.

   The amount of data copied doesn't matter for this demonstration. A small file is enough to demonstrate the restore process.
2. In the **Navigation** pane of the AWS Storage Gateway console, select the gateway.
3. In the **Volumes** tab, select the storage volume created for the gateway.

   This gateway should have only one storage volume. Selecting the volume displays its properties.
4. Click the **Create Snapshot** button to create a snapshot of the volume.

   Depending on the amount of data on the disk and the upload bandwidth, it might take a few seconds to complete the snapshot. Note the volume ID for the volume from which you create a snapshot. You will use the ID to find the snapshot.

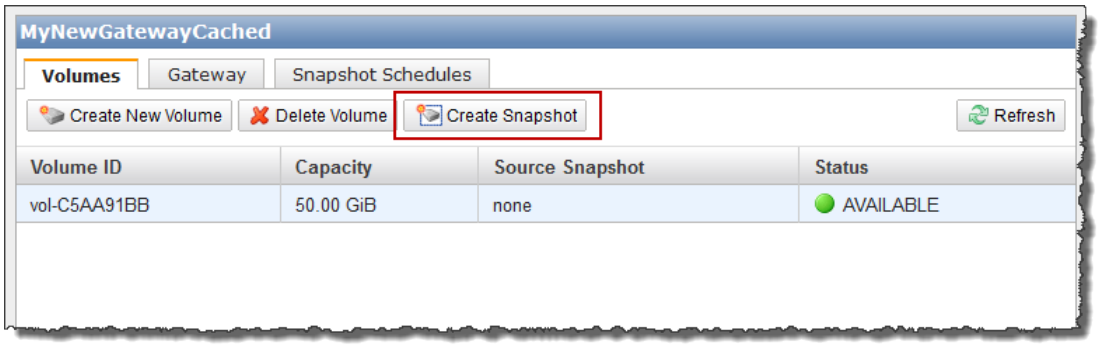


5. In the **Snapshot Taken** confirmation window, click **Close**.

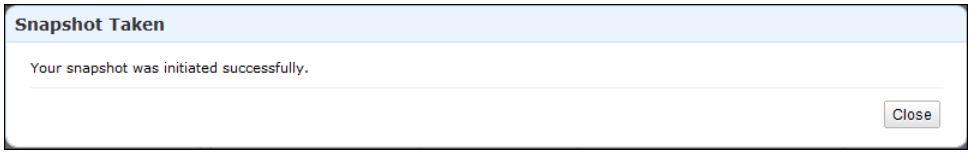


6. In the **Navigation** pane, click **Snapshots**, and find the snapshot that you just created.

   You can use the **Started on** column value and the volume ID you noted earlier to confirm the snapshot's source. Note that the **Started on** time is UTC time.

   The **Status** of your snapshot might be **pending**. In this case, you must wait for the snapshot **Status** to change to **completed** before restoring the snapshot.

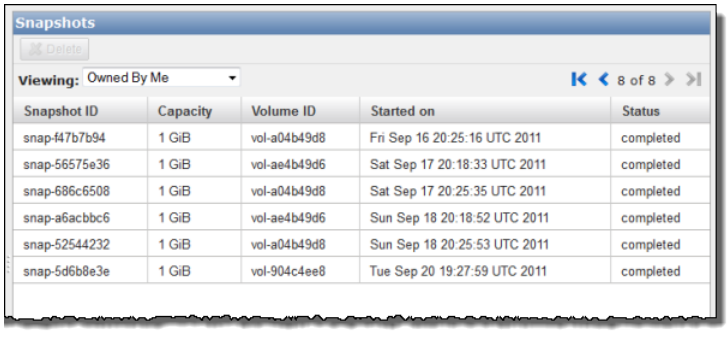
Snapshots
Delete
Viewing: Owned By Me
8 of 8
Snapshot ID
Capacity
Volume ID
Started on
Status
snap-f47b7b94 1 GiB vol-a04b49d8 Fri Sep 16 20:25:16 UTC 2011 completed
snap-56575e36 1 GiB vol-ae4b49d6 Sat Sep 17 20:18:33 UTC 2011 completed
snap-686c6508 1 GiB vol-a04b49d8 Sat Sep 17 20:25:35 UTC 2011 completed
snap-a6acbbc6 1 GiB vol-ae4b49d6 Sun Sep 18 20:18:52 UTC 2011 completed
snap-52544232 1 GiB vol-a04b49d8 Sun Sep 18 20:25:53 UTC 2011 completed
snap-5d6b8e3e 1 GiB vol-904c4ee8 Tue Sep 20 19:27:59 UTC 2011 completed

7. Copy the snapshot ID from the **Snapshot ID** list so you can enter it in a subsequent step when you create a storage volume based on the snapshot.

### To restore the snapshot

1. In the AWS Storage Gateway console, click the name of gateway in the navigation pane.
2. Click the **Volumes** tab, and then click **Create New Volume**.
3. In the **Create Storage Volume** dialog box, enter the following information.

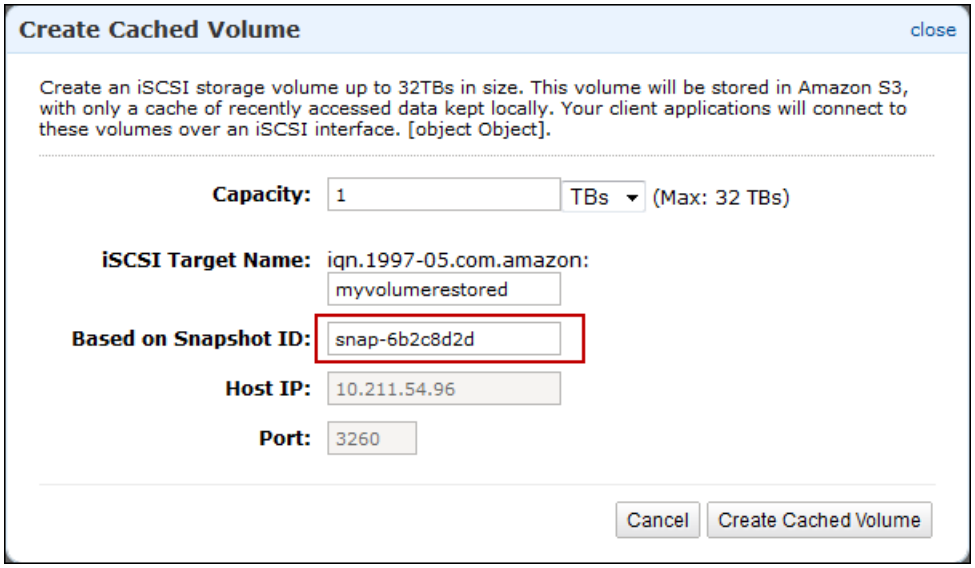


   a. In the **Capacity** text box, enter the same capacity as the original volume from which you took the snapshot.
   b. In the **iSCSI Target Name** box, type a name for your iSCSI target—for example, **`myvolumerestored`**.

      The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** GUI after discovery. For example, the name `target1` appears as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN.
   c. In the **Based on Snapshot ID** field, enter the snapshot ID.
   d. Click **Create Cached Volume**.

      Doing this creates a storage volume based on your snapshot. The volume details appear in the AWS Storage Gateway console.

4. Connect to the new volume target.

   a. On the **Start** menu of your Windows client computer, click **Run**, type **`iscsicpl.exe`**, and then click **OK** to run the iSCSI initiator program.
   b. In the **iSCSI Initiator Properties** dialog box, click the **Targets** tab. If the new target does not appear in the **Discovered Targets** pane, click **Refresh**.

      You should see both the original target and the new target. The new target will have a status of **Inactive**.

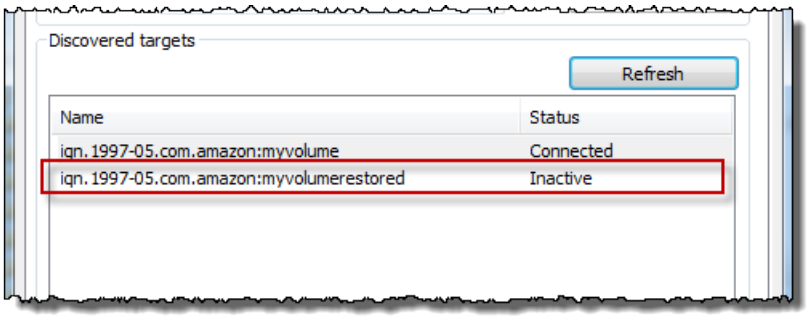


c. Select the new target, and click **Connect**.

d. In the **Connect to Target** dialog box, click **OK**.

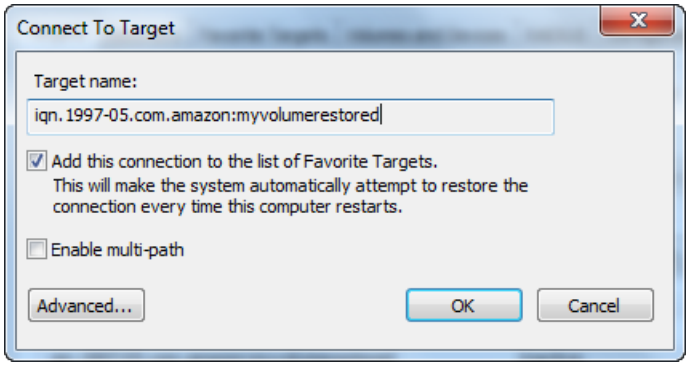
Connect To Target
Target name:
iqn.1997-05.com.amazon:myvolumerestored
Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.
Enable multi-path
Advanced...
OK
Cancel

5. Bring the restored volume online.

   a. If the **Disk Management** console is not already open, on the **Start** menu of your Windows client computer, click **Run**, type **`diskmgmt.msc`**, and then click **OK** to open the console.

      The restored storage volume is shown in the console with a warning.

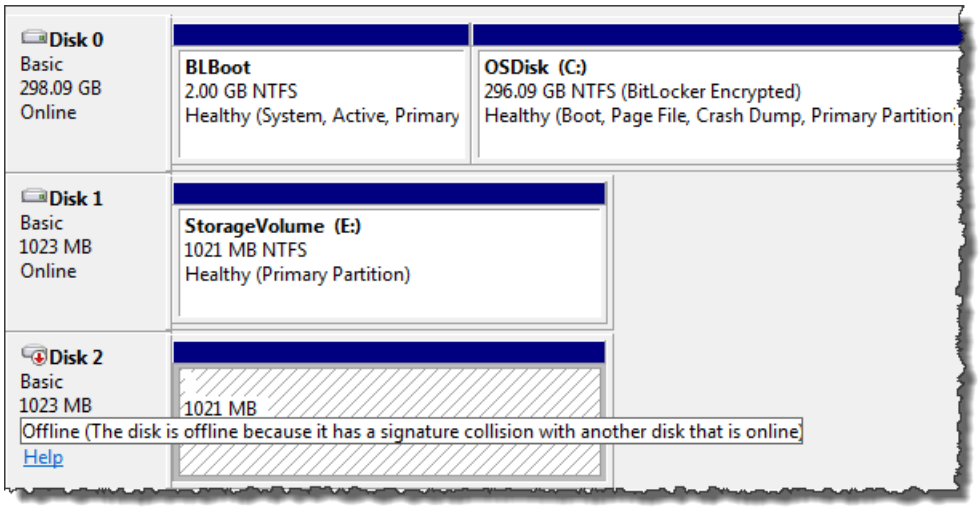


   b. Right-click the restored volume and select **Online**. Doing this brings the volume online and assigns it a different drive letter.

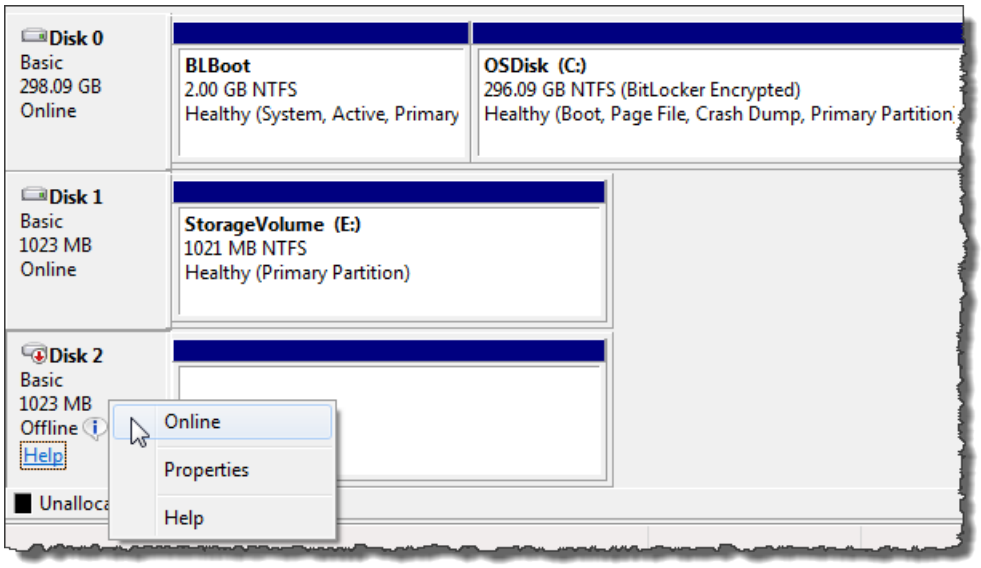


6. Open the restored volume and verify that the data you saved earlier is there.

**Next Step**

Step 6: Clean Up After the Example Exercise (p. 94)

# Test the Setup for a Gateway-Stored Volume

You verify the setup for a gateway-stored volume by taking a snapshot backup of your storage volume and storing the snapshot in AWS. You then restore the snapshot on another storage volume.

> **Note**
> Restoring data from Amazon EBS volumes that are encrypted is not supported.

### To create a snapshot of a storage volume

1. On your Windows computer, copy some data to your mapped storage volume.

   The amount of data copied doesn't matter for this demonstration. A small file is enough to demonstrate the restore process.
2. In the AWS Storage Gateway console, select the gateway in the navigation pane.
3. In the **Volumes** tab, select the storage volume created for the gateway.

   This gateway should have only one storage volume. Selecting the volume displays its properties.
4. Click the **Create Snapshot** button to create a snapshot of the volume.

   Depending on the amount of data on the disk and the upload bandwidth, it might take a few seconds to complete the snapshot. Note the volume ID from which you create a snapshot. You will use the ID to find the snapshot.

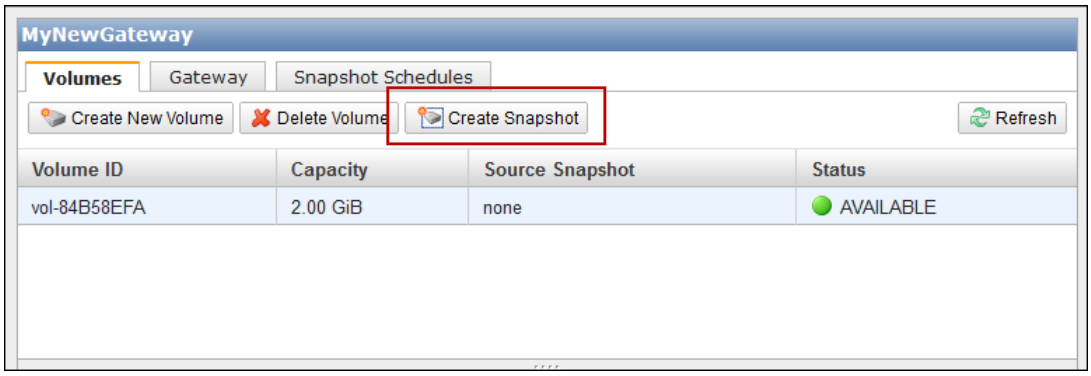


5. In the **Snapshot Taken** confirmation window, click **Close**.

6. In the navigation pane, click **Snapshots**, and find the snapshot that you just created.

   You can use the **Started on** column value and the volume ID you noted earlier to confirm the snapshot's source. Note the **Started on** time is UTC time.

   The **Status** of your snapshot might be **pending**. In this case, you must wait for the snapshot **Status** to change to **completed** before restoring the snapshot.

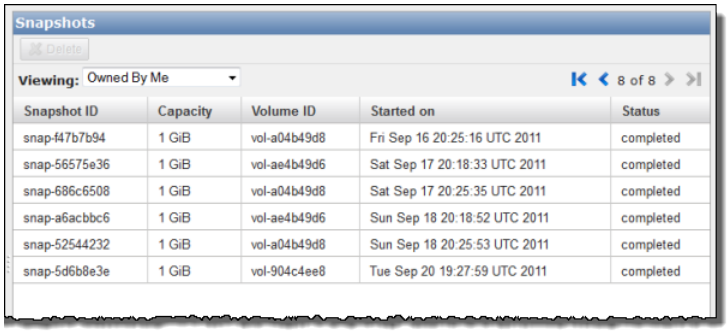
Snapshots
Delete
Viewing: Owned By Me
8 of 8
Snapshot ID
Capacity
Volume ID
Started on
Status
snap-f47b7b94 1 GiB vol-a04b49d8 Fri Sep 16 20:25:16 UTC 2011 completed
snap-56575e36 1 GiB vol-ae4b49d6 Sat Sep 17 20:18:33 UTC 2011 completed
snap-686c6508 1 GiB vol-a04b49d8 Sat Sep 17 20:25:35 UTC 2011 completed
snap-a6acbbc6 1 GiB vol-ae4b49d6 Sun Sep 18 20:18:52 UTC 2011 completed
snap-52544232 1 GiB vol-a04b49d8 Sun Sep 18 20:25:53 UTC 2011 completed
snap-5d6b8e3e 1 GiB vol-904c4ee8 Tue Sep 20 19:27:59 UTC 2011 completed

7. Copy the snapshot ID from the **Snapshot ID** list so you can enter it in a subsequent step when you create a storage volume based on the snapshot.

### To restore the snapshot

1. Add another virtual disk to your VM. This disk will become a new storage volume, on which you will restore the snapshot. Because the original storage volume was 2 GiB in size, create a new virtual disk of the same size. For instructions, see Step 2.3.2: Provision Local Disk Storage (p. 24) in the Getting Started section.
2. In the AWS Storage Gateway console, click the name of the gateway in the navigation pane.
3. Click the **Volumes** tab, and then click **Create New Volume**.
4. In the **Create Storage Volume** dialog box, enter the following information.

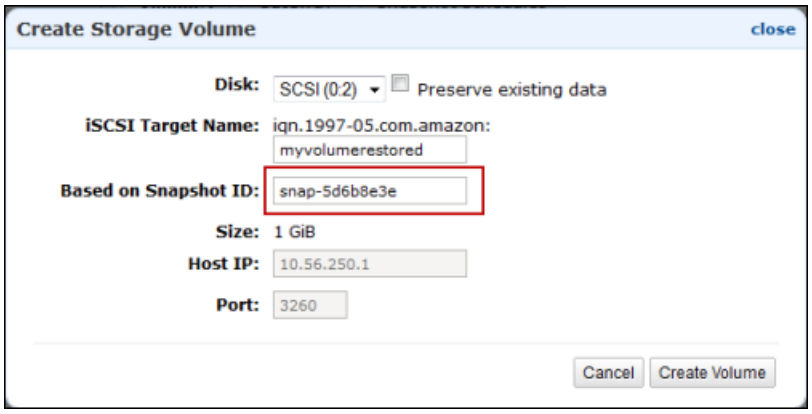


a. In the **Disk** drop-down list, select the virtual disk that you added in the preceding step.
b. In the **iSCSI Target Name** box, enter a name for your iSCSI target—for example, **`myvolumerestored`**.

   The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** GUI after discovery. For example, the name `target1` appears as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN.
c. In the **Based on Snapshot ID** box, type the snapshot ID.
d. Click **Create Volume**.

   Doing this creates a storage volume based on your snapshot. The storage volume details appear in the AWS Storage Gateway console.

5. Connect to the new volume target.

   a. On the **Start** menu of your Windows client computer, click **Run**, type **`iscsicpl.exe`**, and then click **OK** to run the iSCSI initiator program.
   b. In the **iSCSI Initiator Properties** dialog box, click the **Targets** tab. If the new target does not appear in the **Discovered Targets** pane, click **Refresh**.

      You should see both the original target and the new target. The new target will have a status of **Inactive**.

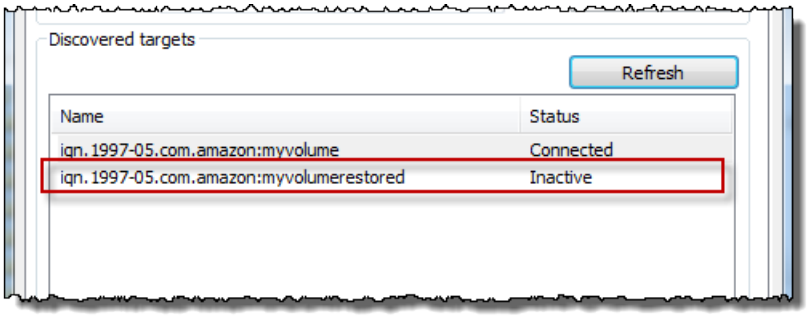


   c. Select the new target, and click **Connect**.
   d. In the **Connect to Target** dialog box, click **OK**.

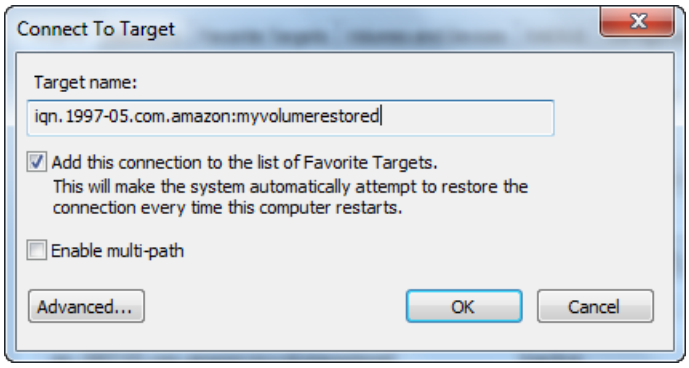
Connect To Target
Target name:
iqn.1997-05.com.amazon:myvolumerestored
Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.
Enable multi-path
Advanced...
OK
Cancel

6. Bring the restored volume online.

   a. If the **Disk Management** console is not already open, on the **Start** menu of your Windows client computer, click **Run**, type **`diskmgmt.msc`**, and then click **OK** to open the console.

      The restored storage volume is shown in the console with a warning.

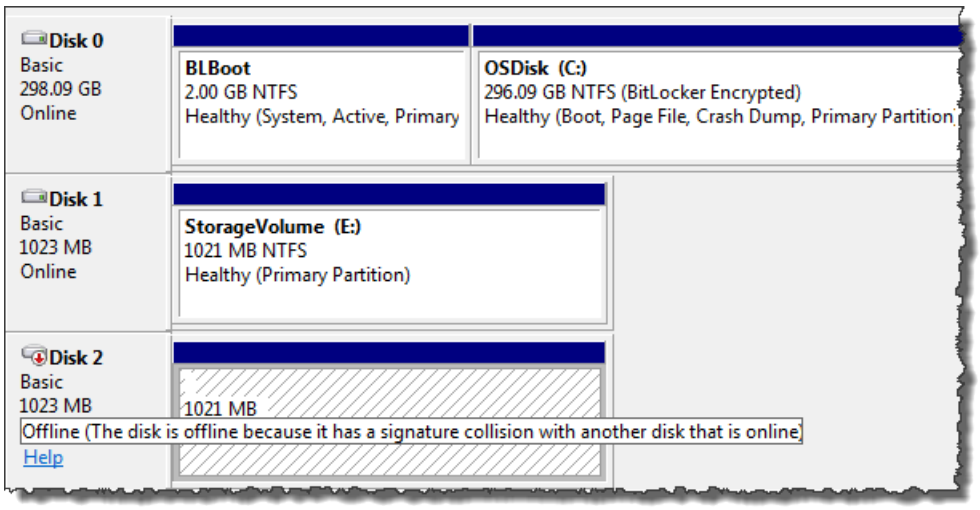


   b. Right-click the restored volume, and select **Online**. Doing this brings the volume online and assigns it a different drive letter.

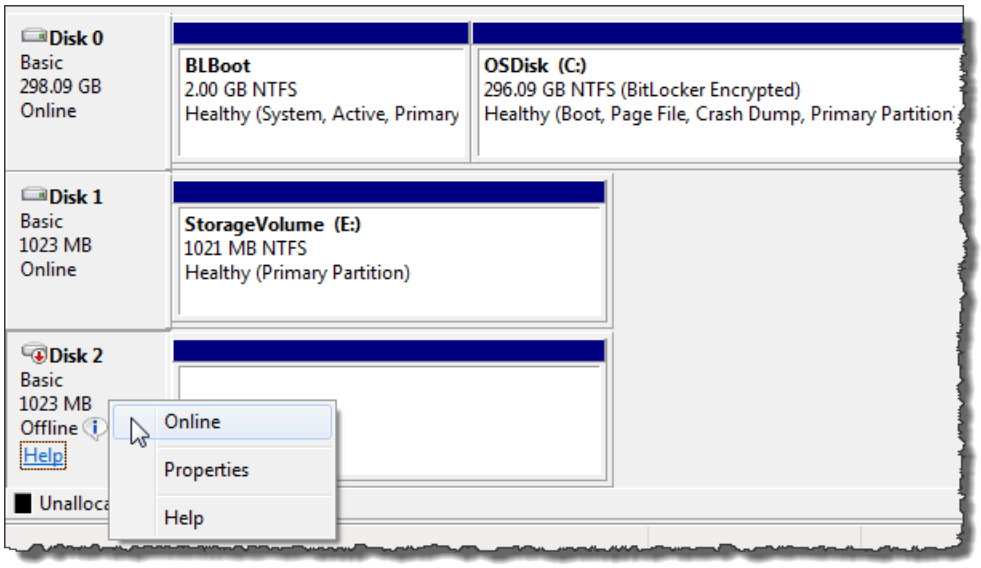


7. Open the restored volume and verify that the data you saved earlier is there.

**Next Step**

Step 6: Clean Up After the Example Exercise (p. 94)

# Step 6: Clean Up After the Example Exercise

If you created the gateway as example exercise, consider cleaning up if you don't plan to continue using your gateway. Doing this avoids incurring charges.

If you plan to continue using your gateway, see additional information in Where Do I Go from Here? (p. 94)

### To free gateway resources

1. Delete any snapshots from the gateway. For instructions, see Deleting a Snapshot (p. 212).
2. You can keep the gateway if you plan to use it. Otherwise, delete the gateway.
   a. After deleting all snapshots, delete the gateway. For instructions, see Deleting a Gateway Using the AWS Storage Gateway Console (p. 136).
   b. Delete the AWS Storage Gateway VM from your on-premises host.

**Next Step**

Where Do I Go from Here? (p. 94)

# Where Do I Go from Here?

The AWS Storage Gateway service provides an easy way for you to back your application storage with the storage infrastructure of the AWS cloud. In Setting Up Your Volume Gateway (Gateway-Cached and Gateway-Stored) (p. 67), you created and provisioned a gateway, and then connected your Windows host to the gateway's storage volume. You added data to the gateway's iSCSI volume, took a snapshot

of the volume and restored it to a new volume, and connected to the new volume and verified that the data shows up on it.

After you finish the exercise, consider the following:

- If you plan on continuing to use your gateway, you should read about sizing the upload buffer more appropriately for real-world workloads. For more information, see Sizing Your Gateway's Storage for Real-World Workloads (p. 95).
- If you don't plan on continuing to use your gateway, consider deleting the gateway to avoid incurring any charges. For more information, see Step 6: Clean Up After the Example Exercise (p. 94).

Other sections of this guide include information about how to do the following:

- Learn more about storage volumes and how to manage them (see Managing Your Activated Gateway (Gateway-Cached and Gateway-Stored) (p. 96)).
- Troubleshoot gateway problems (see Troubleshooting Your Volume Gateway (Gateway-Cached and Gateway-Stored) (p. 226)).
- Optimize your gateway (see Optimizing Gateway Performance (p. 158)).
- Understand Storage Gateway metrics and how you can monitor how your gateway performs (see Monitoring Your Gateway (p. 185)).
- Connect to the gateway's iSCSI targets to store data (see Connecting to Storage Volumes (p. 111)).

## Sizing Your Gateway's Storage for Real-World Workloads

By this point, you have a simple, working gateway. However, the assumptions used to create this gateway are not appropriate for real-world workloads. If you want to use this gateway for real-world workloads, you need to do two things:

1. Size your upload buffer appropriately.
2. Set up monitoring for your upload buffer, if you haven't done so already.

This topic shows how to do both of these tasks. If you activated a gateway for gateway-cached volumes, you also need to size your cache storage for real-world workloads.

### To size your upload buffer and cache storage for a gateway-cached setup

- Use the formula shown in Managing the Upload Buffer and Cache Storage (Gateway-Cached) (p. 169) for sizing the upload buffer. We strongly recommend that you allocate at least 150 GiB for the upload buffer. If the upload buffer formula yields a value less than 150 GiB, use 150 GiB as your allocated upload buffer.

  The upload buffer formula takes into account the difference between throughput from your application to your gateway and throughput from your gateway to AWS, multiplied by how long you expect to write data. For example, assume that your applications write text data to your gateway at a rate of 40 MB per second for 12 hours a day and your network throughput is 12 MB per second. Assuming a compression factor of 2:1 for the text data, the formula specifies that you need to allocate approximately 675 GiB of upload buffer space.

### To size your upload buffer for a gateway-stored setup

- Use the formula discussed in Managing the Upload Buffer (Gateway-Stored) (p. 178). We strongly recommend that you allocate at least 150 GiB for your upload buffer. If the upload buffer formula yields a value less than 150 GiB, use 150 GiB as your allocated upload buffer.

The upload buffer formula takes into account the difference between throughput from your application to your gateway and throughput from your gateway to AWS, multiplied by how long you expect to write data. For example, assume that your applications write text data to your gateway at a rate of 40 MB per second for 12 hours a day and your network throughput is 12 MB per second. Assuming a compression factor of 2:1 for the text data, the formula specifies that you need to allocate approximately 675 GiB of upload buffer space.

### To monitor your upload buffer

1. View your gateway's current upload buffer.

   - In the **Gateway** tab in the AWS Storage Gateway console, find the **Upload Buffer Used** field.

     The following example shows the upload buffer at 50 percent.

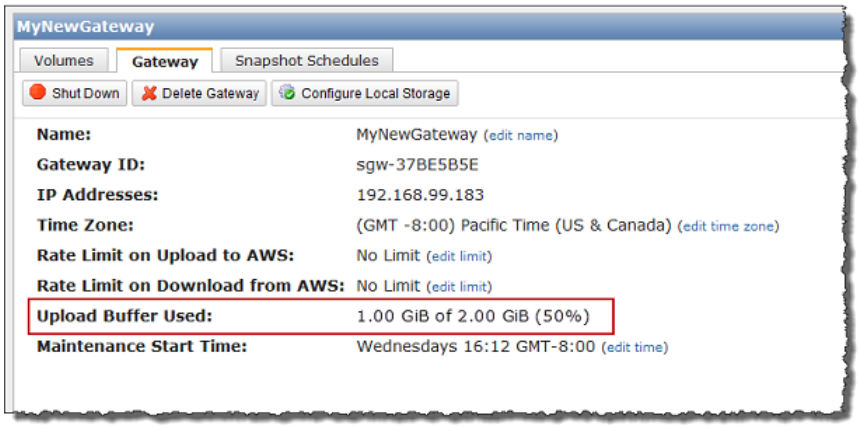


2. Set one or more alarms to notify you about upload buffer use.

   We highly recommend that you create one or more upload buffer alarms in the CloudWatch console. For example, you can set an alarm for a level of use you want to be warned about and an alarm for a level of use that, if exceeded, is cause for action. The action might be adding more upload buffer space. For more information, see To set an upper threshold alarm for a gateway's upload buffer (p. 192).

# Managing Your Activated Gateway (Gateway-Cached and Gateway-Stored)

This section discusses how you can manage your gateway after you have deployed and activated it. Managing your gateway includes tasks such as configuring cache storage and upload buffer space, working with volumes, doing general maintenance, troubleshooting, and monitoring your gateway's performance. If you have not set up a gateway, see Setting Up Your Volume Gateway (Gateway-Cached and Gateway-Stored) (p. 67).

**Topics**

- Managing Storage Volumes in AWS Storage Gateway (p. 97)
- Connecting to Storage Volumes (p. 111)
- Performing Volume Gateway Maintenance Tasks (p. 125)
- Monitoring Your Gateway (p. 185)

- Logging Volume Gateway API Calls by Using AWS CloudTrail (p. 200)
- Working with Snapshots (p. 201)
- Troubleshooting Your Volume Gateway (Gateway-Cached and Gateway-Stored) (p. 226)

# Managing Storage Volumes in AWS Storage Gateway

This section discusses how to manage your existing storage volumes, including adding new volumes, removing existing volumes, and viewing the status of a volume.

**Topics**

- Managing Storage Volumes (Gateway-Cached) (p. 97)
- Managing Storage Volumes (Gateway-Stored) (p. 100)
- Understanding Storage Volume Status (p. 105)

## Managing Storage Volumes (Gateway-Cached)

Gateway-cached volumes are storage volumes in Amazon S3 that are exposed as iSCSI targets on which you can store your application data. This topic discuses how to add and remove volumes from your gateway-cached setup.

**Topics**

- Adding a Storage Volume (p. 97)
- Removing a Storage Volume (p. 99)

> **Important**
> If a cached volume keeps your primary data in Amazon S3, you should avoid processes that read or write all data on the entire volume. For example, we don't recommend using virus scanning software that scans the entire cached volume. Such a scan, whether on-demand or scheduled, will cause all data stored in Amazon S3 to be downloaded locally for scanning, which results in high bandwidth usage. Instead of doing a full disk scan, you can use real-time virus scanning—that is, scanning data as it is read from or written to the cached volume.

Resizing a storage volume is not supported. To change the size of a volume, create a snapshot of the volume, and then create a new cached volume from the snapshot. The new volume can be bigger than the volume from which the snapshot was created. For steps describing how to remove a storage volume, see To remove a storage volume (p. 100). For steps describing how to add a storage volume and preserve existing data, see To create a storage volume using the console (p. 98).

All gateway-cached volume data and snapshot data is stored in Amazon S3 and is encrypted at rest using Server Side Encryption (SSE). However, you cannot access this data using the Amazon S3 API or with other tools such as the Amazon S3 console.

### Adding a Storage Volume

As your application needs grow, you might need to add more storage volumes to your gateway. As you add more storage volumes, you must consider the size of the cache storage and upload buffer you allocated to the gateway. The gateway must have sufficient buffer and cache space for new volumes. For more information, see Managing the Upload Buffer and Cache Storage (Gateway-Cached) (p. 169).

You can add storage volumes using the AWS Storage Gateway console or AWS Storage Gateway API. For information on using the AWS Storage Gateway API to add storage volumes, see

CreateCachediSCSIVolume. The following procedure demonstrates using the console and assumes that you already have a deployed and activated gateway.

### To create a storage volume using the console

1. In the **Navigation** pane of the AWS Storage Gateway console, select the gateway. If the gateway is activated but has no storage volumes, see Step 2: Configure Local Storage and Alarms (Gateway-Cached and Gateway-Stored) (p. 67) in the Getting Started section for instructions on how to specify cache storage and upload buffer storage and then create volumes.
2. If the gateway already has one or more storage volumes, perform the following steps.

   On the **Volumes** tab, click **Create New Volume**.

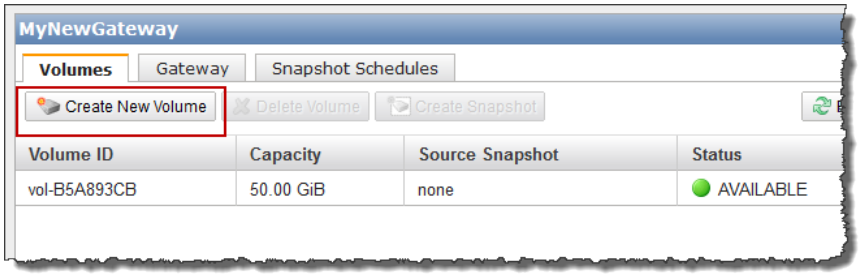


3. In the Configure Your Activated Gateway wizard, configure the volume.

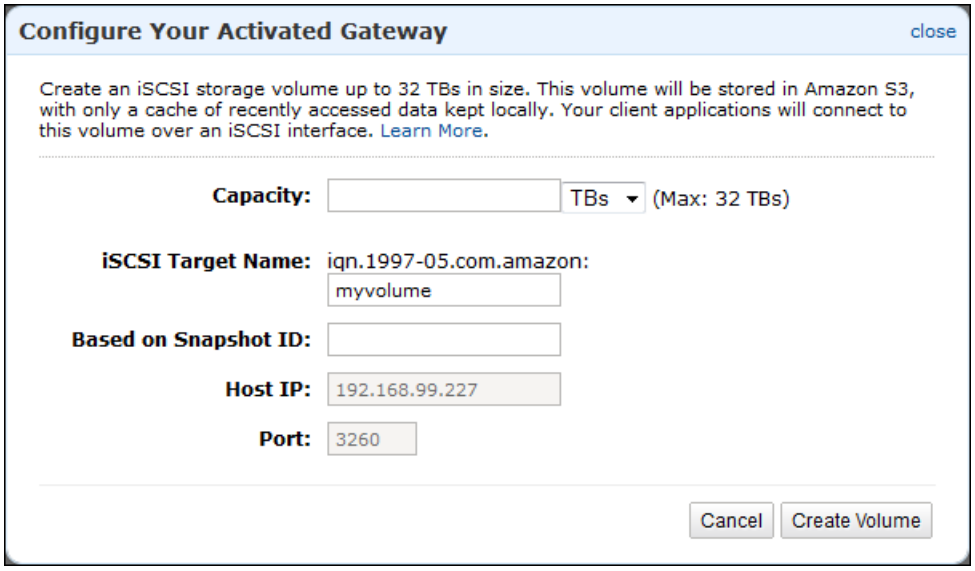
Configure Your Activated Gateway
close
Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3, with only a cache of recently accessed data kept locally. Your client applications will connect to this volume over an iSCSI interface. Learn More.
Capacity:
TBs
(Max: 32 TBs)
iSCSI Target Name:
iqn.1997-05.com.amazon:
myvolume
Based on Snapshot ID:
Host IP:
192.168.99.227
Port:
3260
Cancel
Create Volume

   a. In the **Capacity** box, type the size of the storage volume to create.

      **Note**
      Resizing a storage volume is not supported. To change the size of a volume, create a snapshot of the volume, and then create a new cached volume from the snapshot. The new volume can be bigger than the volume from which the snapshot was created. For steps describing how to remove a storage volume, see To remove a storage volume (p. 100). For steps describing how to add a storage volume and preserve existing data, see To create a storage volume using the console (p. 98).

   b. In the size drop-down list next to the **Capacity** box, select the appropriate size of the volume, GBs or TBs. For information size limits and other limits, see Appendix F: AWS Storage Gateway Limits (p. 421)

c. Type a name in the **iSCSI Target Name** box.

   The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name on the **Targets** tab of the **iSCSI Microsoft Initiator** UI after discovery. For example, the name `target1` appears as `iqn.1997-05.com.amazon:target1`. Ensure that the target name is globally unique within your storage area network (SAN).

d. If you are creating a volume from a snapshot, type the snapshot ID in the **Based on Snapshot ID** box.

   You can specify the ID of an existing AWS Storage Gateway or Amazon EBS snapshot you previously created. In this case, the gateway creates the storage volume and downloads data to the local cache storage only on first access of the data. For information about how to find a snapshot you want to use, see Finding a Snapshot (p. 203).

   **Important**
   If the Amazon EBS volume you create the snapshot from contains data from a public Amazon Machine Image (AMI), such as an AMI from AWS Marketplace, you will not have the permission to create the volume.
   Storage volumes cannot be created from snapshots that are encrypted.
   The size of the storage volume you create must be equal to or greater than the size of the snapshot you create it from. For information on how to add a disk to your gateway VM that can be used as a gateway-stored volume, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23). Once you add such a volume, you can access the contents of the volume from your on-premises applications. For more information, see Connecting to Storage Volumes (p. 111).

e. If you've configured your local gateway host with multiple network interface cards (NICs), specify the IP address for the NIC to use for this storage volume in the **Host IP** box.

f. Check that the **Port** box shows 3260. The **Port** box shows the port to which to map an iSCSI target. AWS Storage Gateway supports only port 3260.

g. Click **Create Volume**.

   Doing this creates a storage volume and makes your disk available as an iSCSI target for your applications to connect to and store data on. For information on connecting to the iSCSI target, see Connecting to Storage Volumes (p. 111).

   **Note**
   If you want snapshots for this volume, you can either take specific a one-time snapshot or set up a snapshot schedule for the volume. For more information, see Editing a Snapshot Schedule (p. 210).

## Removing a Storage Volume

You might need to remove a storage volume as your application needs change—for example, if you migrate your application to use a larger storage volume. Before removing a storage volume, make sure that there are no applications currently writing to the volume. Also, make sure that there are no snapshots in progress for the volume. If a snapshot schedule is defined for the volume, you can check it on the **Snapshot Schedules** tab of the AWS Storage Gateway console. For more information, see Editing a Snapshot Schedule (p. 210).

You can remove storage volumes using the AWS Storage Gateway console or the AWS Storage Gateway API. For information on using the AWS Storage Gateway API to remove storage volumes, see DeleteVolume. The following procedure demonstrates using the console.

### To remove a storage volume

1. On the **Volumes** tab of the AWS Storage Gateway console, select the storage volume.

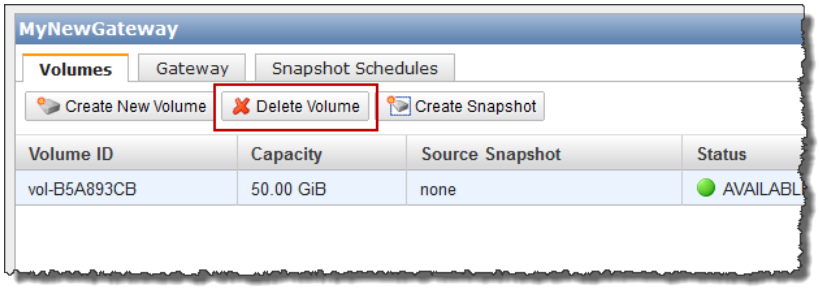


2. Click **Delete Volume**.

# Managing Storage Volumes (Gateway-Stored)

Gateway-stored volumes are storage volumes that are exposed as iSCSI targets on which you can store your application data. The storage volumes are created on the local virtual disks that you added to your gateway VM—for more information on that process, see On-Premises Gateway (p. 15). This topic discuses how to add and remove volumes from your gateway-stored setup.

**Topics**

- Adding a Storage Volume (p. 100)
- Removing a Storage Volume (p. 102)

Resizing the underlying disk of a storage volume is not supported. To change the size of an underlying disk, delete the storage volume that is using the disk, resize the disk, and then create a new storage volume from the resized disk. When you recreate the storage volume, be sure to preserve the data on the disk. For steps describing how to remove a storage volume, see To remove the underlying local disk (VMware ESXi) (p. 103) or To remove the underlying local disk (Microsoft Hyper-V) (p. 104). For steps describing how to add a storage volume and preserve existing data, see To create a storage volume using the console (p. 100).

## Adding a Storage Volume

As your application needs grow, you might need to add more storage volumes to your gateway. As you add more storage volumes, you must consider the size of your upload buffer you allocated to the gateway. The gateway must have sufficient buffer space. For more information, see Managing the Upload Buffer (Gateway-Stored) (p. 178).

You can add storage volumes using the AWS Storage Gateway console or the AWS Storage Gateway API. For information on using the API to add storage volumes, see CreateStorediSCSIVolume. The following procedure demonstrates using the console and assumes that you already have a deployed and activated gateway. Furthermore, the procedure assumes that there is at least one locally provisioned disk of the gateway that is not used and can be allocated as a storage volume. For information on how to provision a local disk for application storage, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23).

### To create a storage volume using the console

1. In the **Navigation** pane of the AWS Storage Gateway console, select the gateway. If the gateway is activated but has no storage volumes, see Step 2: Configure Local Storage and Alarms

(Gateway-Cached and Gateway-Stored) (p. 67) in the Getting Started section for instructions on how to specify the cache storage and upload buffer and then create volumes.

2. Click **Create New Volume**.

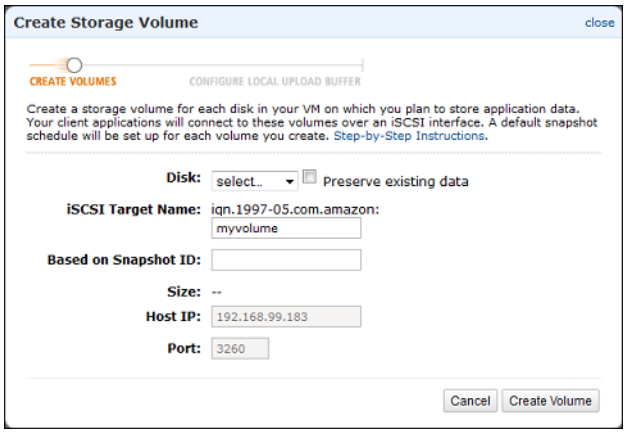
Create Storage Volume
close
CREATE VOLUMES
CONFIGURE LOCAL UPLOAD BUFFER
Create a storage volume for each disk in your VM on which you plan to store application data. Your client applications will connect to these volumes over an iSCSI interface. A default snapshot schedule will be set up for each volume you create. Step-by-Step Instructions.
Disk:
select..
Preserve existing data
iSCSI Target Name:
iqn.1997-05.com.amazon:
myvolume
Based on Snapshot ID:
Size:
--
Host IP:
192.168.99.183
Port:
3260
Cancel
Create Volume

3. In the Create Storage Volume wizard, specify the following information:

   a. In the **Disk** list, select the local virtual disk that you provisioned for the gateway.

      For information about provisioning disks, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23).

   b. If you want to preserve data on the disk, click to select the **Preserve existing data** check box.

      If you preserve data, AWS Storage Gateway bootstraps your volume upon creation, uploading your volume's existing data to AWS. This process might take some time to complete.

      When you create a volume from an existing snapshot and you don't want to preserve any existing data on the disk, make sure the **Preserve existing data** check box is not selected.

   c. Type a name in the **iSCSI Target Name** box.

      The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name on the **Targets** tab of the **iSCSI Microsoft Initiator** UI after discovery. For example, the name `target1` appears as `iqn.1997-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN.

   d. If you are creating a volume from a snapshot, specify the snapshot ID in the **Based on Snapshot ID** box.

      You can specify the ID of an existing AWS Storage Gateway or Amazon EBS snapshot that you previously created. This approach is useful if you want to restore a snapshot of another storage volume. In this case, the gateway creates the storage volume and downloads your existing snapshot data to the volume. However, you don't need to wait for all of the data to transfer from Amazon S3 to your volume before your application can start accessing the volume and all of its data. For more information about snapshots, see Working with Snapshots (p. 201).

      **Important**
      If the Amazon EBS volume you create the snapshot from contains data from a public Amazon Machine Image (AMI), such as an AMI from AWS Marketplace, you will not have the permission to create the volume.
      Storage volumes cannot be created from snapshots that are encrypted.
      The size of the storage volume you create must be equal to or greater than the size of the snapshot you create it from. For information on how to add a disk to your gateway VM that can be used as a gateway-stored volume, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23). Once you add such a volume, you can

access the contents of the volume from your on-premises applications. For more information, see Connecting to Storage Volumes (p. 111).

By default, AWS limits the maximum size of Amazon EBS volumes you can create per AWS account. For information about default Amazon EBS volume limits, see Amazon EBS Limits. in the *Amazon EC2 User Guide*. For information about how to increase these limits, go to Amazon EC2 Service Limits.

e. If you've configured your local gateway host with multiple network interface cards (NICs), specify the IP address for the NIC to use for this storage volume in the **Host IP** box.

f. Check that the **Port** box shows 3260. The **Port** box shows the port to which to map an iSCSI target. AWS Storage Gateway supports only port 3260.

g. Click **Create Volume**.

Doing this creates a storage volume and makes your disk available as an iSCSI target for your applications to connect to and store data on.

Clicking this button also creates a snapshot schedule for your new volume. By default, AWS Storage Gateway takes snapshots once a day. You can change both the time the snapshot occurs each day and the frequency (every 1, 2, 4, 8, 12, or 24 hours). For more information, see Editing a Snapshot Schedule (p. 210).

**Note**
Snapshots are incremental, compressed backups. For a given storage volume, the gateway saves only the blocks that have changed since the last snapshot. This approach minimizes the amount of storage that is used for your backups.

## Removing a Storage Volume

You might need to remove a storage volume as your application needs change—for example, if you migrate your application to use a larger storage volume and you want to reclaim the underlying, local disk space of the old storage volume. To reclaim the local disk space, you need to remove the local disk from the VM.

Before removing a storage volume, make sure that there are no applications currently writing to the volume. Also, make sure that there are no snapshots in progress for the volume. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see Editing a Snapshot Schedule (p. 210).

You can remove storage volumes using the AWS Storage Gateway console or the AWS Storage Gateway API. For information on using the API to remove storage volumes, see DeleteVolume. The following procedure demonstrates using the console and either the vSphere client for a gateway deployed on the VMware ESXi platform or the Microsoft Hyper-V Manager for a gateway deployed on the Microsoft Hyper-V platform.

### To remove a storage volume on the AWS Storage Gateway console

**Note**
Perform the following procedure with the gateway running.

1. In the **Volumes** tab of the AWS Storage Gateway console, select the storage volume.
2. If you plan to remove the disk from the VM that backs the storage volume, click the **Details** tab, and note the value in the **Local Disk** box.

   This value is the disk's virtual device node value, which you use in the hypervisor client to help ensure that you remove the correct disk.

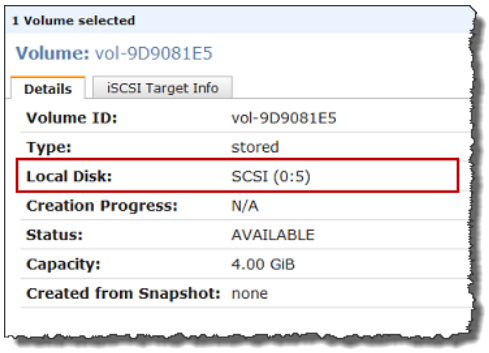


3. Click **Delete Volume**.
4. If you also want to remove the underlying local disk, do one of the following:

| For a Gateway Hosted In... | Do This... |
|---|---|
| VMware ESXi | Follow the steps in To remove the underlying local disk (VMware ESXi) (p. 103). |
| Microsoft Hyper-V | Follow the steps in To remove the underlying local disk (Microsoft Hyper-V) (p. 104). |

### To remove the underlying local disk (VMware ESXi)

1. In the vSphere client, right-click the name of your gateway VM and click **Edit Settings**.
2. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk to remove, and then click **Remove**.

   Verify that the **Virtual Device Node** value in the **Virtual Machine Properties** dialog box has the same value that you noted in step 2 of the preceding procedure. Doing this helps ensure that you remove the correct disk. The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.

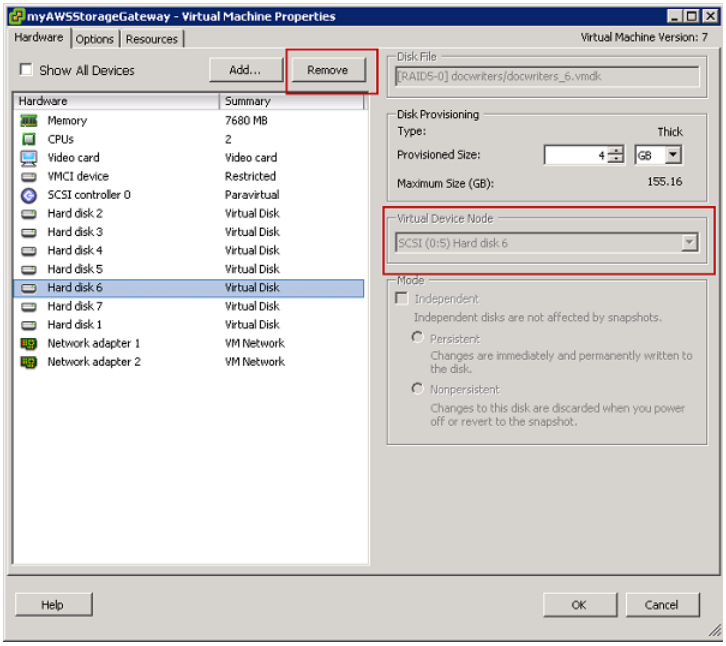


3. Choose the appropriate option in the **Removal Options** panel, and then click **OK** to complete the process of removing the disk.

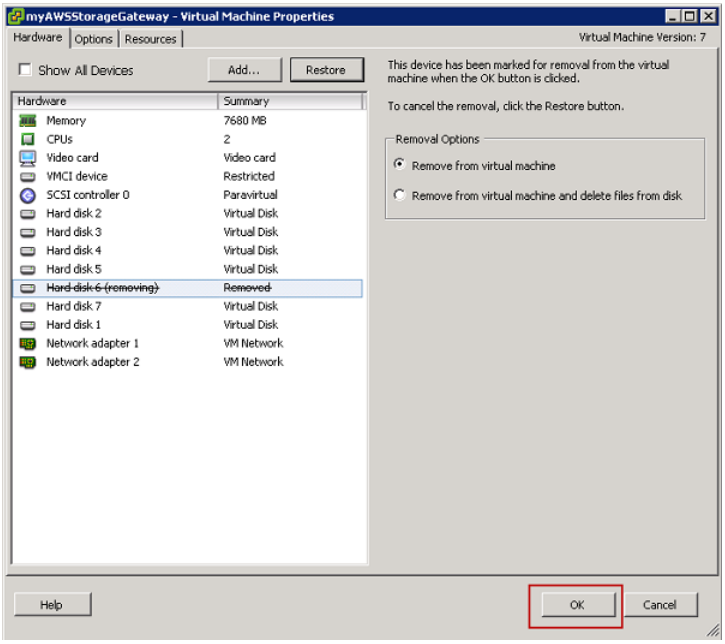


### To remove the underlying local disk (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, right-click the name of your gateway VM and click **Settings**.
2. In the **Hardware** list of the **Settings** dialog box, select the disk to remove, and then click **Remove**.

   The disks you add to a gateway are under the **SCSI Controller** entry in the **Hardware** list.

   Verify that the **Controller** and **Location** value are the same value that you noted from a previous step. Doing this helps ensure that you remove the correct disk.

The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.

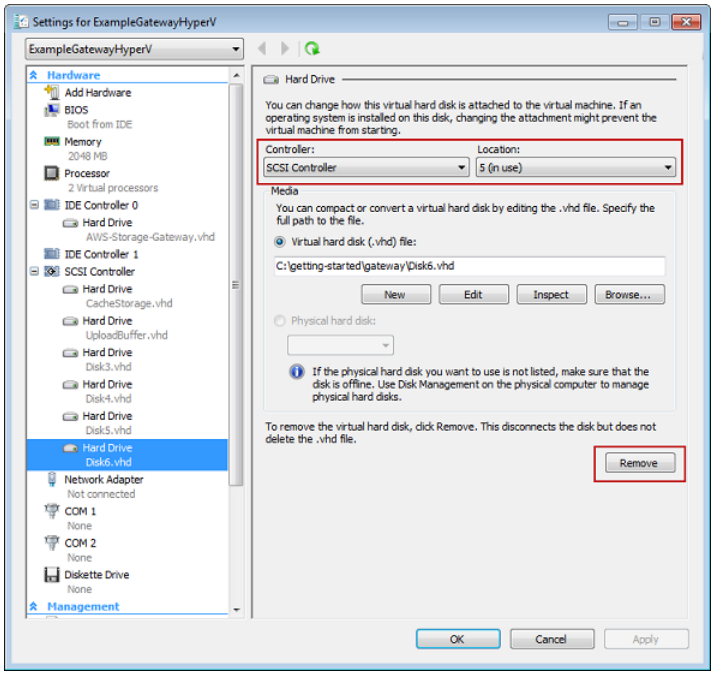


3. Click **OK** to apply the change.

## Understanding Storage Volume Status

Each storage volume has an associated status that tells you at a glance what the health of the volume is. Most of the time, the status will indicate that the volume is functioning normally and that no action is needed on your part. In a few cases, the status indicates a problem with the volume that might or might not require action on your part. This section helps you decide when you need to take action.

You can see volume status on the AWS Storage Gateway console or by using one of the AWS Storage Gateway API operations, for example DescribeCachediSCSIVolumes or DescribeStorediSCSIVolumes. The following example shows volume status on the AWS Storage Gateway console. Volume status appears in the **Status** column for each storage volume on your gateway. In the example, you can tell that the volume highlighted is functioning normally because the status is AVAILABLE.

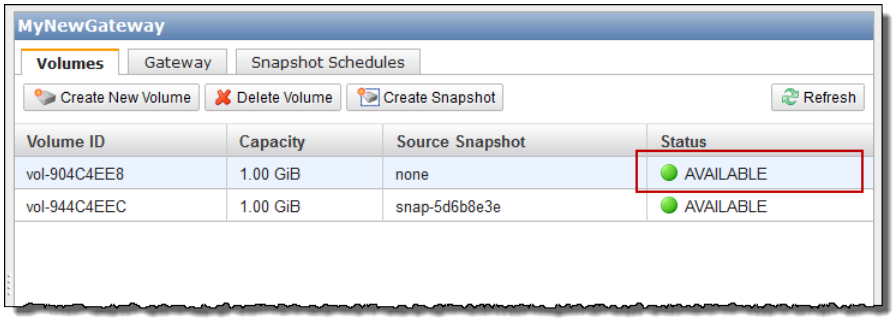


The following table describes each storage volume status, and if and when you should take action based on each status. The AVAILABLE status is the normal status of a volume, and a volume should have this status all or most of the time it is in use.

| Status | Meaning |
| --- | --- |
| AVAILABLE | The volume is available for use. The AVAILABLE status is the normal running status for a volume. |
| BOOTSTRAPPING | The gateway is synchronizing data locally with a copy of the data stored in AWS. You typically do not need to take any action for this status, because the storage volume will automatically go to the AVAILABLE status in most cases.<br><br>The following are three scenarios when a volume status is BOOTSTRAPPING.<br><br>• A gateway has unexpectedly shut down.<br>• A gateway's upload buffer has been exceeded. In this scenario, bootstrapping occurs when your volume has the PASS THROUGH status and the amount of free upload buffer increases sufficiently. You can provide additional upload buffer space as one way to increase the percentage of free upload buffer space. In this particular scenario, the storage volume goes from PASS THROUGH to BOOTSTRAPPING to AVAILABLE status. You can continue to use this volume during this bootstrapping period. However, you cannot take snapshots of the volume.<br>• You are creating a gateway-stored volume and preserving existing local disk data. In this scenario, your gateway starts uploading all of the data up to AWS and the volume has the BOOTSTRAPPING status until all of the data from the local disk is copied to AWS. You can use the volume during this bootstrapping period. However, you cannot take snapshots of the volume. |
| CREATING | The volume is currently being created and is not ready to be used. The CREATING status is transitional. No action is required. |
| DELETING | The volume is currently being deleted. The DELETING status is transitional. No action is required. |
| IRRECOVERABLE | An error occurred from which the volume cannot recover. For information on taking action in this situation, see Troubleshooting Storage Volume Issues (p. 233). |

| Status | Meaning |
| --- | --- |
| PASS THROUGH | Data maintained locally is out of sync with data stored in AWS. Note that data written to a volume while the volume is in PASS THROUGH status remains in the cache until the volume status is BOOTSTRAPPING, and starts to upload to AWS when BOOTSTRAPPING status begins.<br><br>The PASS THROUGH status can occur for several reasons.<br><br>• The PASS THROUGH status occurs if your gateway has run out of upload buffer space. Your applications can continue to read from and write data to your storage volumes while the volumes have the PASS THROUGH status. However, the gateway is not writing any of your volume data to its upload buffer or uploading any of this data to AWS. The gateway will continue to upload any data written to the volume before the volume entered the PASS THROUGH status. Any pending or scheduled snapshots of a storage volume will fail while the volume has the PASS THROUGH status. For information about what action to take when your storage volume has the PASS THROUGH status because the upload buffer has been exceeded, see Troubleshooting Storage Volume Issues (p. 233).<br>• The PASS THROUGH status occurs when there is more than one storage volume bootstrapping at once. Only one gateway storage volume can bootstrap at a time. For example, if you create two storage volumes and choose to preserve existing data on both of them, then the second storage volume will have the PASS THROUGH status until the first storage volume finishes bootstrapping. In this scenario, you do not need to take action. Each storage volume will change to the AVAILABLE status automatically when it is finished being created. You can read and write to the storage volume while it has the PASS THROUGH or BOOTSTRAPPING status.<br>• Infrequently, the PASS THROUGH status can indicate that a disk allocated for upload buffer use has failed. For information about what action to take in this scenario, see Troubleshooting Storage Volume Issues (p. 233). |
| RESTORING | The volume is being restored from an existing snapshot. This status applies only for gateway-stored volumes. For more information, see How AWS Storage Gateway Works (Architecture) (p. 2).<br><br>If you restore two storage volumes at the same time, both storage volumes will show RESTORING as their status. Each storage volume will change to the AVAILABLE status automatically when it is finished being created. You can read and write to a storage volume and take a snapshot of it while it has the RESTORING status. |

| Status | Meaning |
| --- | --- |
| RESTORING PASS THROUGH | The volume is being restored from an existing snapshot and has encountered an upload buffer issue. This status applies only for gateway-stored volumes. For more information, see How AWS Storage Gateway Works (Architecture) (p. 2).<br><br>One reason that can cause the RESTORING PASS THROUGH status is if your gateway has run out of upload buffer space. Your applications can continue to read from and write data to your storage volumes while they have the RESTORING PASS THROUGH status. However, no snapshots of a storage volume can occur during the RESTORING PASS THROUGH status period. For information about what action to take when your storage volume has the RESTORING PASS THROUGH status because upload buffer capacity has been exceeded, see Troubleshooting Storage Volume Issues (p. 233).<br><br>Infrequently, the RESTORING PASS THROUGH status can indicate that a disk allocated for an upload buffer has failed. For information about what action to take in this scenario, see Troubleshooting Storage Volume Issues (p. 233). |
| UPLOAD BUFFER NOT CONFIGURED | The volume cannot be created or used, because the gateway does not have an upload buffer configured. For information on how to add upload buffer capacity for volumes in a gateway-cached setup, see Adding and Removing Upload Buffer Capacity (Gateway-Cached) (p. 173). For information on how to add upload buffer capacity for volumes in a gateway-stored setup, see Adding and Removing Upload Buffer Capacity (Gateway-Stored) (p. 180). |

## Cached Volume Status Transitions (Gateway-Cached)

The following state diagram describes the most common transitions between gateway-cached volume statuses. It is not necessary for you to understand the diagram in detail to use your gateway effectively. Rather, the diagram provides detailed information if you are interested in understanding more about how AWS Storage Gateway works.

The diagram shows neither the UPLOAD BUFFER NOT CONFIGURED status nor the DELETING status. Volume states in the diagram are represented by green, yellow, and red boxes. The colors are interpreted as follows.

Green
: The gateway is operating normally. The volume status is AVAILABLE or will eventually become AVAILABLE.

Yellow
: The volume has the PASS THROUGH status, which indicates there is a potential issue with the storage volume. If this status is because the upload buffer space is filled, then in some cases buffer space will become available again. At that point, the storage volume self-corrects to the AVAILABLE status. In other cases, you might have to add more upload buffer space to your gateway to allow the storage volume status to become AVAILABLE. For information on how to troubleshoot a case when upload buffer capacity has been exceeded, see Troubleshooting Storage Volume Issues (p. 233). For information on how to add upload buffer capacity, see Adding and Removing Upload Buffer Capacity (Gateway-Cached) (p. 173).

Red
: The storage volume has the IRRECOVERABLE status. In this case, you should delete the volume. For information on how to do this, see To remove a storage volume (p. 100).

In the diagram, a transition between two states is depicted with a labeled line. For example, the transition from the CREATING status to the AVAILABLE status is labeled as *Create Basic Volume or Create Volume*

*from Snapshot* and represents creating a cached volume. For more information about creating storage volumes, see Adding a Storage Volume (p. 97).

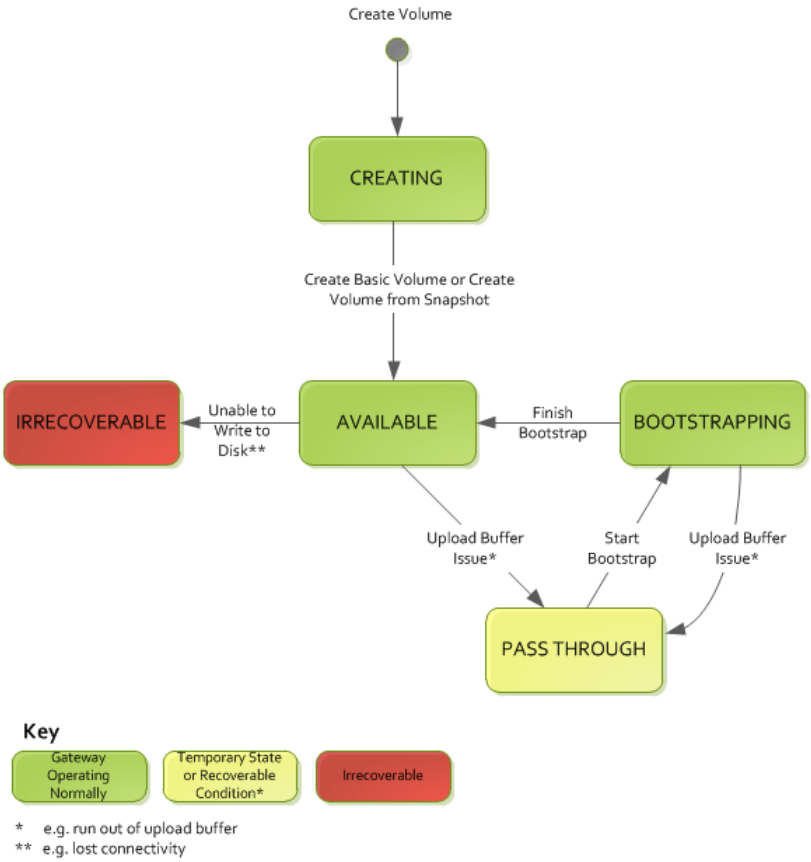


**Note**
The volume status of PASS THROUGH is depicted as yellow in this diagram and does not match the color of this status icon in the **Status** box of the AWS Storage Gateway console.

## Stored Volume Status Transitions (Gateway-Stored)

The following state diagram describes the most common transitions between gateway-stored volume statuses. It is not necessary for you to understand the diagram in detail to use your gateway effectively. Rather, the diagram provides detailed information if you are interested in understanding more about how AWS Storage Gateway works.

The diagram shows neither the UPLOAD BUFFER NOT CONFIGURED status nor the DELETING status. Volume states in the diagram are represented by green, yellow, and red boxes. The colors are interpreted as follows.

Green
: The gateway is operating normally. The volume status is AVAILABLE or will eventually become AVAILABLE.

Yellow
: When you are creating a storage volume and preserving data, then the path from the CREATING status to the PASS THROUGH status occurs if another volume is bootstrapping. In this case, the volume with the PASS THROUGH status will go to the BOOTSTRAPPING status and then to the AVAILABLE status when the first volume is finished bootstrapping. Other than the specific scenario mentioned, yellow (PASS THROUGH status) indicates that there is a potential issue with the storage volume, the most common one being an upload buffer issue. If upload buffer capacity has been

exceeded, then in some cases buffer space will become available again. At that point, the storage volume self-corrects to the AVAILABLE status. In other cases, you might have to add more upload buffer capacity to your gateway to return the storage volume to the AVAILABLE status. For information on how to troubleshoot a case when upload buffer capacity has been exceeded, see Troubleshooting Storage Volume Issues (p. 233). For information on how to add upload buffer capacity, see Managing the Upload Buffer (Gateway-Stored) (p. 178).

Red

The storage volume has the IRRECOVERABLE status. In this case, you should delete the volume. For information on how to do this, see To remove the underlying local disk (VMware ESXi) (p. 103) or To remove the underlying local disk (Microsoft Hyper-V) (p. 104).

In the following diagram, a transition between two states is depicted with a labeled line. For example, the transition from the CREATING status to the AVAILABLE status is labeled as *Create Basic Volume* and represents creating a storage volume without preserving data or creating the volume from a snapshot. For more information about creating a storage volume, see To create a storage volume using the console (p. 100).

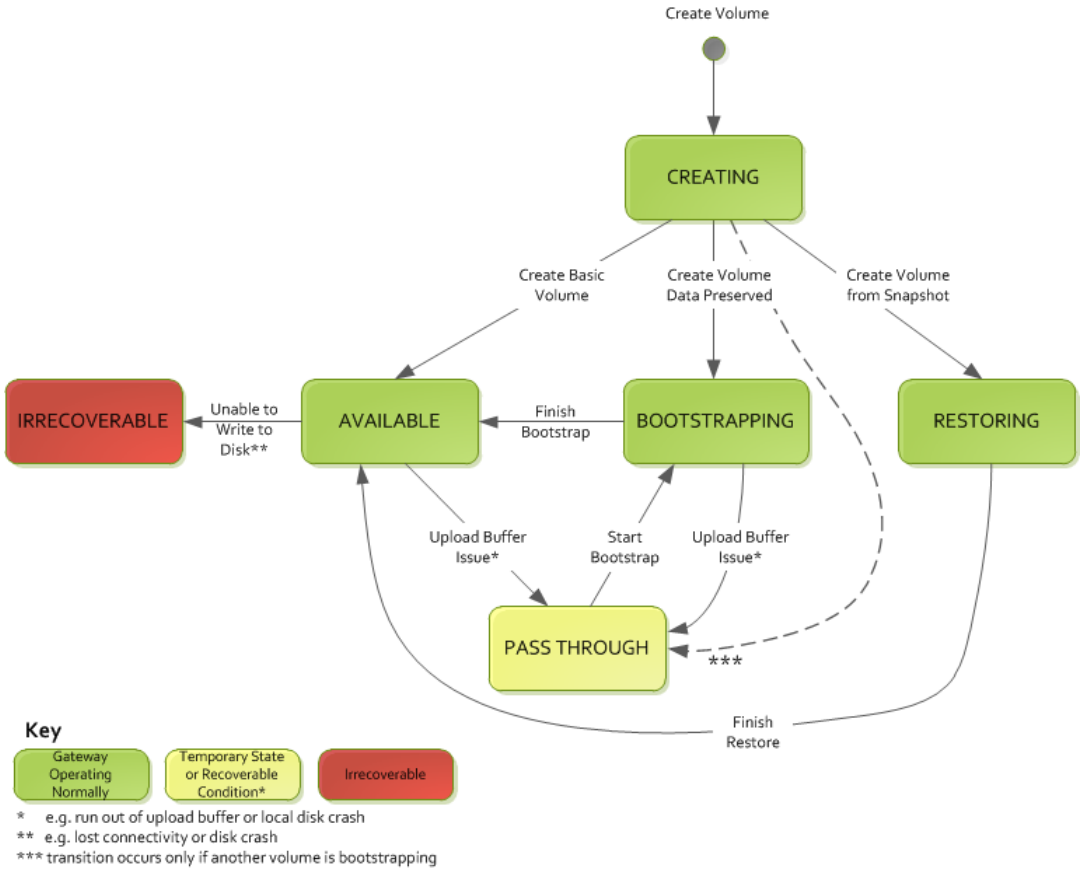


**Note**

The volume status of PASS THROUGH is depicted as yellow in this diagram and does not match the color of this status icon in the **Status** box of the AWS Storage Gateway console.

# Connecting to Storage Volumes

The gateway exposes volumes you have created for the gateway as Internet Small Computer System Interface (iSCSI) targets. For more information, see Step 4: Connect Your Volumes to Your Windows Client (p. 79).

> **Note**
> To connect to your storage volume target, your gateway must have an upload buffer configured. If an upload buffer is not configured for your gateway, then the status of your storage volumes is displayed as UPLOAD BUFFER NOT CONFIGURED. To configure an upload buffer for a gateway in a gateway-stored setup, see To configure additional upload buffer space for your gateway (p. 181). To configure an upload buffer for a gateway in a gateway-cached setup, see Follow the steps in To configure a local disk as an upload buffer for your gateway (p. 184).

**Topics**

- Customizing Your Windows iSCSI Settings (p. 112)
- Connecting from a Red Hat Client to Your Storage Volume (p. 114)
- Configuring CHAP Authentication for Your Storage Volume (p. 116)

The following diagram highlights the iSCSI target in the larger picture of the AWS Storage Gateway architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 2).

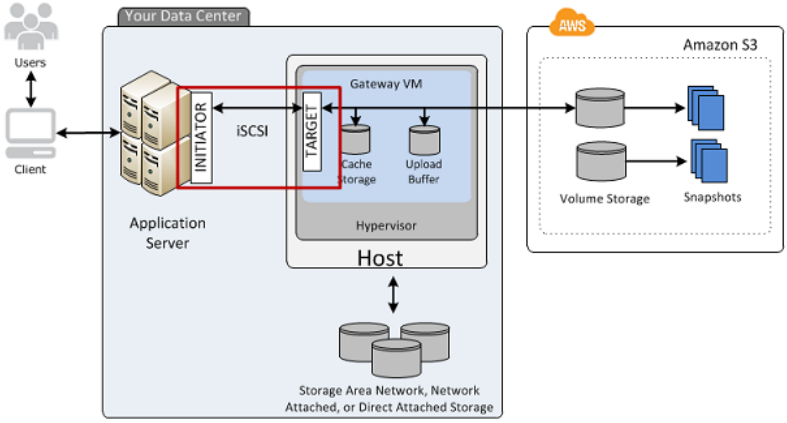


The following list defines some of the iSCSI terms that are used to describe the connection and the components involved.

**iSCSI initiator**
The client component of an iSCSI network. The initiator sends requests to the iSCSI target. Initiators can be implemented in software or hardware. AWS Storage Gateway only supports software initiators.

**iSCSI target**
The server component of the iSCSI network that receives and responds to requests from initiators. Each of your storage volumes is exposed as an iSCSI target. Connect only one iSCSI initiator to each iSCSI target.

**Microsoft iSCSI initiator**
The software program on Windows computers that enables you to connect a client computer (that is, the computer running the application whose data you want to write to the gateway) to an external iSCSI-based array (that is, the gateway). The connection is made using the host computer's Ethernet network adapter card. The Microsoft iSCSI initiator is already installed on Windows Server 2008 R2, Windows 7, Windows Server 2008, and Windows Vista. On these operating systems, you do not need to install the initiator.

**Red Hat iSCSI initiator**

The iscsi-initiator-utils Resource Package Manager (RPM) package provides you with an iSCSI initiator implemented in software for Red Hat. The package includes a server daemon for the iSCSI protocol.

You can connect to your storage volume from either a Windows or Red Hat client. You can optionally configure Challenge-Handshake Authentication Protocol (CHAP) for either client type.

Your gateway exposes your volume as an iSCSI target with a name you specify, prepended by `iqn.1997-05.com.amazon:`. For example, if you specify a target name of `myvolume`, then the iSCSI target you use to connect to the volume is `iqn.1997-05.com.amazon:myvolume`. For more information about how to configure your applications to mount volumes over iSCSI, see Connecting to Storage Volumes (p. 111).

| To... | See... |
|---|---|
| Connect to your storage volume from Windows. | Step 4: Connect Your Volumes to Your Windows Client (p. 79) in the Getting Started section |
| Connect to your storage volume from Red Hat Linux. | Connecting from a Red Hat Client to Your Storage Volume (p. 114) |
| Configure CHAP authentication for Windows and Red Hat Linux. | Configuring CHAP Authentication for Your Storage Volume (p. 116) |

# Customizing Your Windows iSCSI Settings

When using a Windows client, you use the Microsoft iSCSI initiator to connect to your gateway storage volume. For instructions on how to connect to your storage volumes, see Step 4: Connect Your Volumes to Your Windows Client (p. 79).

After setting up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets. By increasing the iSCSI timeout values as shown in the following steps, you improve the ability of your application to deal with write operations that take a long time and other transient issues such as network interruptions.

**Note**

Before making changes to the registry, you should make a backup copy of the registry. For information on making a backup copy and other best practices to follow when working with the registry, go to Registry best practices in the *Microsoft TechNet Library*.

### To customize your Windows iSCSI settings

1. Increase the maximum time for which requests are queued.
   a. Start Registry Editor (`Regedit.exe`).
   b. Navigate to the globally unique identifier (GUID) key for the device class that contains iSCSI controller settings, shown following.

      **Warning**

      Make sure you are working in the **CurrentControlSet** subkey and not another control set such as **ControlSet001** or **ControlSet002**.

```
HKEY_Local_Machine\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-
11CE-BFC1-08002BE10318}
```

c. Find the subkey for the Microsoft iSCSI initiator, shown following as *<Instance Number>*.

The key will be represented by a four-digit number such as `0000` or `00001`.

```
HKEY_Local_Machine\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>
```

Depending on what is installed on your computer, the Microsoft iSCSI initiator might not be the subkey `0000`. You can ensure that you have selected the correct subkey by verifying that the string `DriverDesc` has the value `Microsoft iSCSI Initiator` as shown in the following example.

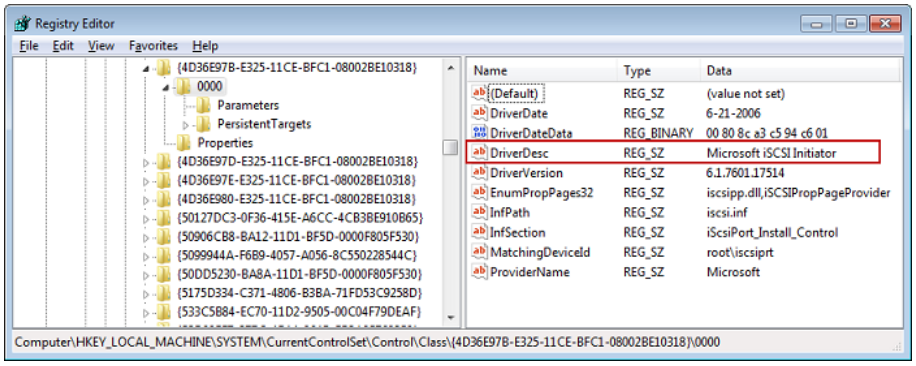


d. Click the **Parameters** subkey to show the iSCSI settings.

e. Right-click the **MaxRequestHoldTime** DWORD (32-bit) value, right-click **Modify**, and change the value to 600.

This value represents a hold time of 600 seconds. The example following shows the **MaxRequestHoldTime** DWORD value with a value of 600.

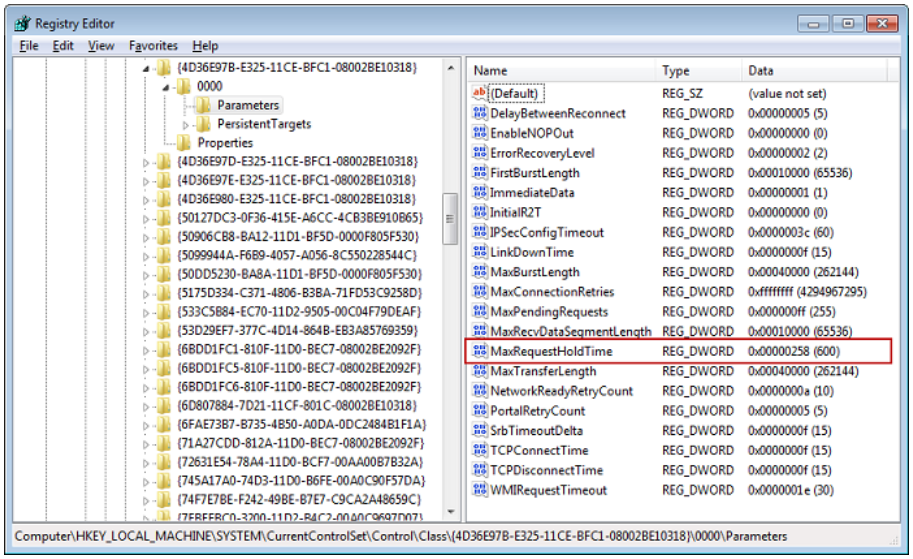


2. Increase the disk timeout value.

a. Start Registry Editor (`Regedit.exe`).

b. Navigate to the **Disk** subkey in the **Services** subkey of the **CurrentControlSet**, shown following.

```
HKEY_Local_Machine\SYSTEM\CurrentControlSet\Services\Disk
```

c. Right-click the **TimeoutValue** DWORD (32-bit) value, click **Modify**, and change the value to 600.

This value represents a timeout of 600 seconds.

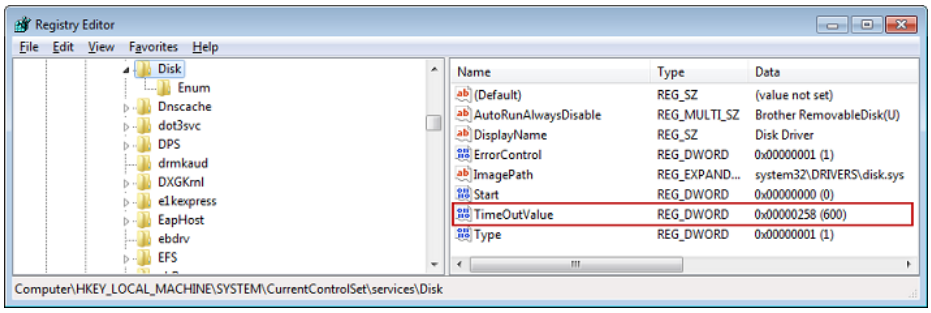


3. Restart your system to ensure that the new configuration values take effect.

Before restarting, you must make sure that the results of all write operations to storage volumes are flushed. To do this, take any mapped storage volume disks offline before restarting.

## Connecting from a Red Hat Client to Your Storage Volume

When using Red Hat Linux, you use the iscsi-initiator-utils RPM package to connect to your gateway storage volume.

### To connect a Linux client to the storage volume

1. Install the iscsi-initiator-utils RPM package, if it isn't already installed on your client.

You can use the following command to install the package.

```
sudo yum install iscsi-initiator-utils
```

2. Ensure that the iSCSI daemon is running.

a. Verify that the iSCSI daemon is running using the following command.

```
sudo /etc/init.d/iscsi status
```

b. If the status command does not return a status of *running*, then start the daemon using the following command.

```
sudo /etc/init.d/iscsi start
```

3. To discover the storage volume targets defined for a gateway, use the following discovery command.

```
sudo /sbin/iscsiadm --mode discovery --type sendtargets --portal GATE
WAY_IP:3260
```

Substitute your gateway's IP address for the *GATEWAY_IP* variable in the preceding command. You can find the gateway IP in the **iSCSI Target Info** properties of a volume on the AWS Storage Gateway console.

The output of the discovery command will look like the following example output.

```
GATEWAY_IP:3260, 1 iqn.1997-05.com.amazon:myvolume
```

Your iSCSI qualified name (IQN) will be different than what is shown preceding, because IQN values are unique to an organization. The name of the target is the name that you specified when you created the storage volume. You can also find this target name in the **iSCSI Target Info** properties pane when you select a storage volume on the AWS Storage Gateway console.

4. To connect to a target, use the following command.

   You need to specify the correct *GATEWAY_IP* and IQN in the connect command.

   **Warning**
   For storage gateways that are deployed on an Amazon EC2 instance, accessing the gateway over a public Internet connection is not supported. The elastic IP address of the Amazon EC2 instance cannot be used as the target address.

```
sudo /sbin/iscsiadm --mode node --targetname iqn.1997-05.com.amazon:myvolume
 --portal GATEWAY_IP:3260,1 --login
```

5. To verify that the volume is attached to the client machine (the initiator), use the following command.

```
ls -l /dev/disk/by-path
```

   The output of the command will look like the following example output.

```
lrwxrwxrwx. 1 root root 9 Apr 16 19:31
ip-GATEWAY_IP:3260-iscsi-iqn.1997-05.com.amazon:myvolume-lun-0 -> ../../sda
```

   After setting up your initiator, we highly recommend that you customize your iSCSI settings as discussed in Customizing Your Linux iSCSI Settings (p. 115).

## Customizing Your Linux iSCSI Settings

After setting up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets. By increasing the iSCSI timeout values as shown following, you improve the ability of your application to deal with write operations that take a long time and other transient issues such as network interruptions.

### To customize your Linux iSCSI settings

1. Increase the maximum time for which requests are queued.

   a. Open the `/etc/iscsi/iscsid.conf` file and find the following lines.

```
node.session.timeo.replacement_timeout = [replacement_timeout_value]
node.conn[0].timeo.noop_out_interval = [noop_out_interval_value]
node.conn[0].timeo.noop_out_timeout = [noop_out_timeout_value]
```

b. Set the *[replacement_timeout_value]* value to 600.

   Set the *[noop_out_interval_value]* value to 60.

   Set the *[noop_out_timeout_value]* value to 600.

   All three values are in seconds.

   **Note**
   The `iscsid.conf` settings must be made before discovering the gateway. If you have already discovered your gateway and/or logged in to the target, you can delete the entry from the discovery database using the following command. Then you can rediscover or log in again to pick up the new configuration.

   ```
   iscsiadm -m discoverydb -t sendtargets -p gateway_ip:3260 -o delete
   ```

2. Increase the disk timeout value in the rules file.

   a. If you are using the Red Hat Enterprise Linux 5 initiator, open the `/etc/udev/rules.d/50-udev.rules` file and find the following line.

   ```
   ACTION=="add", SUBSYSTEM=="scsi" , SYSFS{type}=="0|7|14", \
   RUN+="/bin/sh -c 'echo [timeout] > /sys$$DEVPATH/timeout'"
   ```

   **Note**
   This rules file does not exist in Red Hat Enterprise Linux 6 initiators, so you must create it.

   To create the rules file in Red Hat Enterprise Linux 6, use the following command and then add the lines of code as shown above.

   ```
   sudo vim /etc/udev/rules.d/50-udev.rules
   ```

   b. Set the *timeout* value to 600.

   This value represents a timeout of 600 seconds.

3. Restart your system to ensure that the new configuration values take effect.

   Before restarting, you must make sure that the results of all write operations to your storage volumes are flushed. To do this, unmount storage volumes before restarting.

4. You can test the configuration by using the following command.

   ```
   udevadm test PATH_TO_ISCSI_DEVICE
   ```

   This command shows the udev rules that are applied to the iSCSI device.

## Configuring CHAP Authentication for Your Storage Volume

AWS Storage Gateway supports authentication between your gateway and iSCSI initiators via Challenge-Handshake Authentication Protocol (CHAP). CHAP provides protection against playback

attacks by periodically verifying the identity of an iSCSI initiator as authenticated to access a storage volume target. To set up CHAP, you must configure it both on the AWS Storage Gateway console and in the iSCSI initiator software that you use to connect to the target. This section discusses mutual CHAP, which is when the initiator authenticates the target and the target authenticates the initiator.

### To set up mutual CHAP for your storage volume

1. Configure CHAP on the AWS Storage Gateway console, as discussed in To configure CHAP on the AWS Storage Gateway console (p. 117).
2. In your client initiator software, complete the CHAP configuration.

   - To configure mutual CHAP on a Windows client, see To configure mutual CHAP on a Windows client (p. 118).
   - To configure mutual CHAP on a Red Hat Linux client, see To configure mutual CHAP on a Red Hat Linux client (p. 123).

### To configure CHAP on the AWS Storage Gateway console

In this procedure, you specify two secret keys that are used to read and write to a storage volume. These same keys are used in the procedure to configure the client initiator.

1. On the AWS Storage Gateway console, go to the **iSCSI Target Info** tab of the volume for which you want to configure CHAP.
2. Click the **Configure CHAP Authentication** link.

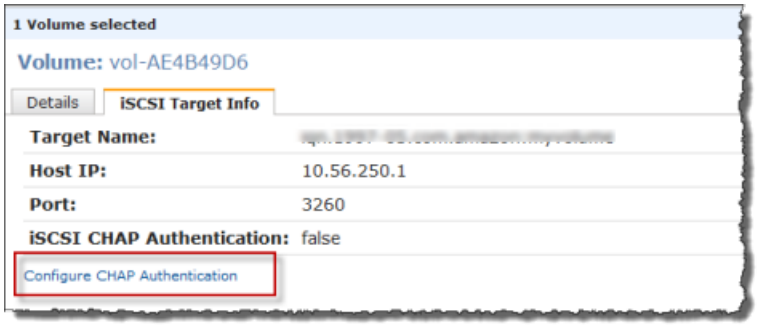


3. Configure CHAP in the **Configure CHAP Authentication** dialog box.

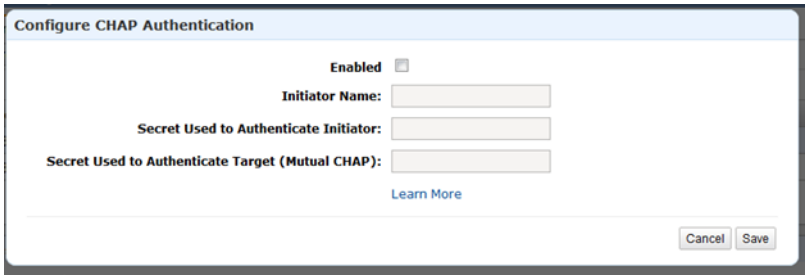


   a. Click to select the **Enabled** check box.
   b. Type the name of your iSCSI initiator in the **Initiator Name** box.

      The initiator name can be found using your iSCSI initiator software. For example, for Windows clients, the name is the value on the **Configuration** tab of the iSCSI initiator. For more information, see To configure mutual CHAP on a Windows client (p. 118).

**Note**
To change an initiator name, you must first disable CHAP, change the initiator name in your iSCSI initiator software, and then enable CHAP with the new name.

c. Type the secret requested in the **Secret used to Authenticate Initiator** box.

This secret must be at least 12 characters long. It is the secret key that the initiator (that is, the Windows client) must know to participate in CHAP with the target.

d. Type the secret requested in the **Secret used to Authenticate Target (Mutual CHAP)** box.

This secret must be at least 12 characters long. It is the secret key that the target must know to participate in CHAP with the initiator.

**Note**
The secret used to authenticate the target must be different than the secret to authenticate the initiator.

e. Click **Save**.

4. On the **iSCSI Target Info** tab of the confirmation dialog box, confirm that CHAP authentication is used.

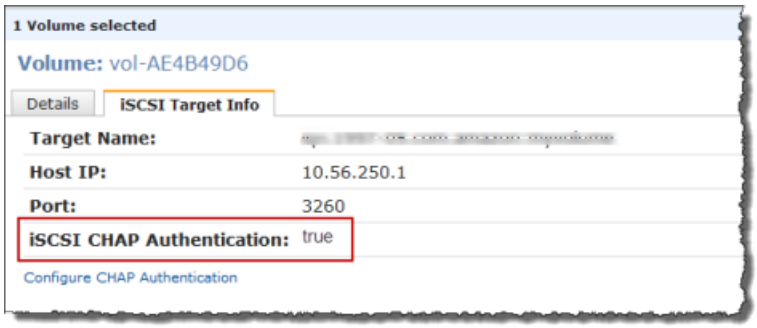


5. Click **Close** in the confirmation dialog box.

### To configure mutual CHAP on a Windows client

In this procedure, you configure CHAP in the Microsoft iSCSI initiator using the same keys that you used to configure CHAP for the storage volume on the console.

1. If the iSCSI initiator is not already started, on the **Start** menu of your Windows client computer, click **Run**, type **`iscsicpl.exe`**, and then click **OK** to run the program.
2. Configure mutual CHAP configuration for the initiator (that is, the Windows client).

   a. Click the **Configuration** tab.

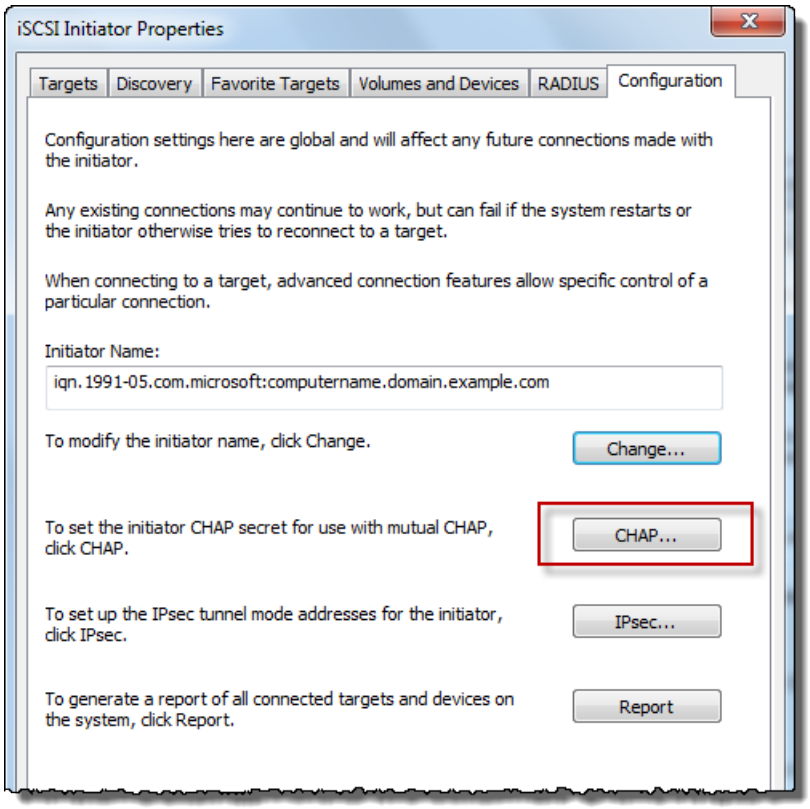
iSCSI Initiator Properties
Targets
Discovery
Favorite Targets
Volumes and Devices
RADIUS
Configuration
Configuration settings here are global and will affect any future connections made with the initiator.
Any existing connections may continue to work, but can fail if the system restarts or the initiator otherwise tries to reconnect to a target.
When connecting to a target, advanced connection features allow specific control of a particular connection.
Initiator Name:
iqn.1991-05.com.microsoft:computername.domain.example.com
To modify the initiator name, click Change.
Change...
To set the initiator CHAP secret for use with mutual CHAP, click CHAP.
CHAP...
To set up the IPsec tunnel mode addresses for the initiator, click IPsec.
IPsec...
To generate a report of all connected targets and devices on the system, click Report.
Report

**Note**

Note that the **Initiator Name** box will be unique to your initiator and company. The name shown here is the value that you used in the **Configure CHAP Authentication** dialog box of the AWS Storage Gateway console.

The name shown in the example image is for demonstration purposes only.

b. Click **CHAP**.

c. In the **iSCSI Initiator Mutual Chap Secret** dialog box, type the mutual CHAP secret value.

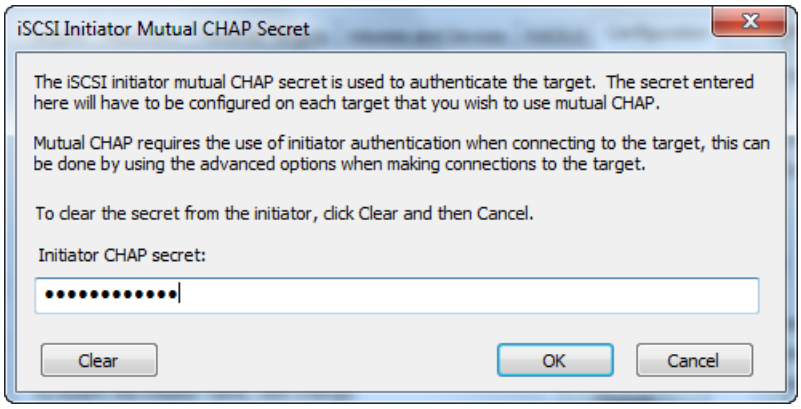
iSCSI Initiator Mutual CHAP Secret
The iSCSI initiator mutual CHAP secret is used to authenticate the target. The secret entered here will have to be configured on each target that you wish to use mutual CHAP.
Mutual CHAP requires the use of initiator authentication when connecting to the target, this can be done by using the advanced options when making connections to the target.
To clear the secret from the initiator, click Clear and then Cancel.
Initiator CHAP secret:
Clear
OK
Cancel

In this dialog box, you are entering the secret that the initiator (the Windows client) uses to authenticate the target (the storage volume). This secret allows the target to read and write to the initiator. This secret maps to the secret typed into the **Secret used to Authenticate Target (Mutual CHAP)** box in the **Configure CHAP Authentication** dialog box. For more information, see Configuring CHAP Authentication for Your Storage Volume (p. 116).

d. If the key that you typed is less than 12 characters, an **Initiator CHAP secret** error dialog box appears.

   Click **OK**, and then type the key again.

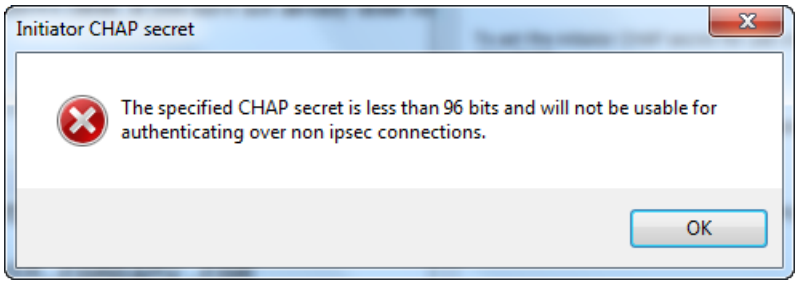
Initiator CHAP secret
The specified CHAP secret is less than 96 bits and will not be usable for authenticating over non ipsec connections.
OK

3. Configure the target with the initiator's secret to complete the mutual CHAP configuration.

   a. Click the **Targets** tab.

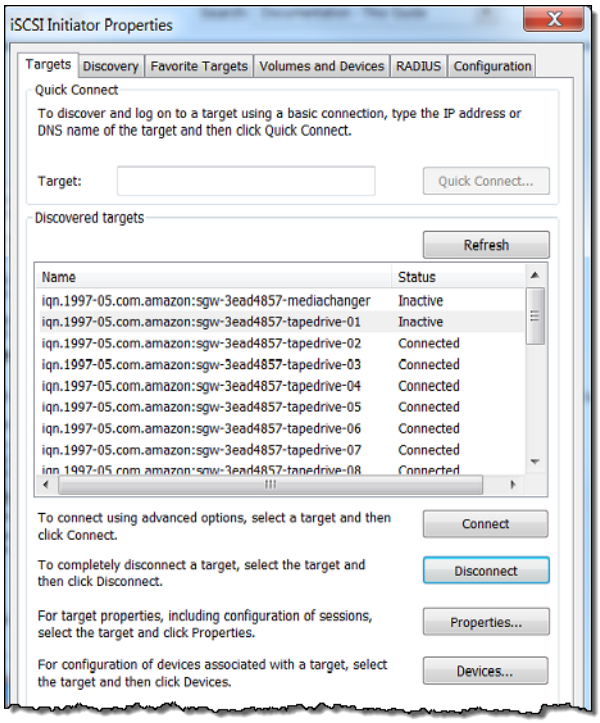


   b. If the target that you want to configure for CHAP is currently connected, disconnect the target by selecting it and clicking **Disconnect**.

   c. Select the target that you want to configure for CHAP, and then click **Connect**.

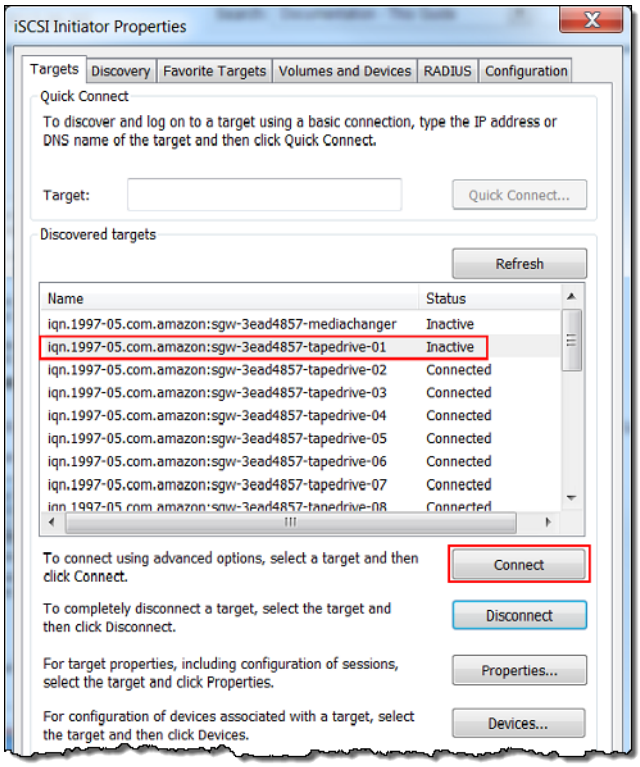


d. In the **Connect to Target** dialog box, click **Advanced**.

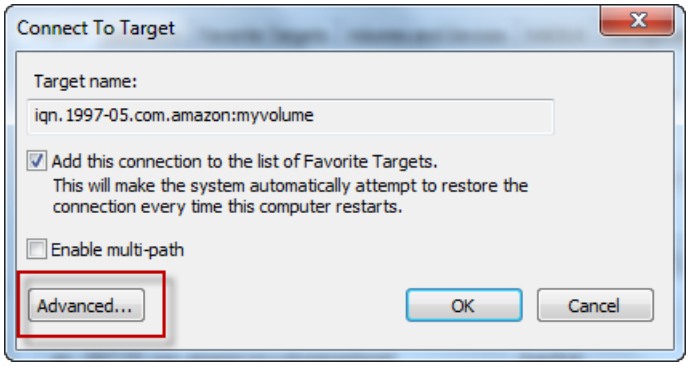
Connect To Target
Target name:
iqn.1997-05.com.amazon:myvolume
Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.
Enable multi-path
Advanced...
OK
Cancel

e. In the **Advanced Settings** dialog box, configure CHAP.

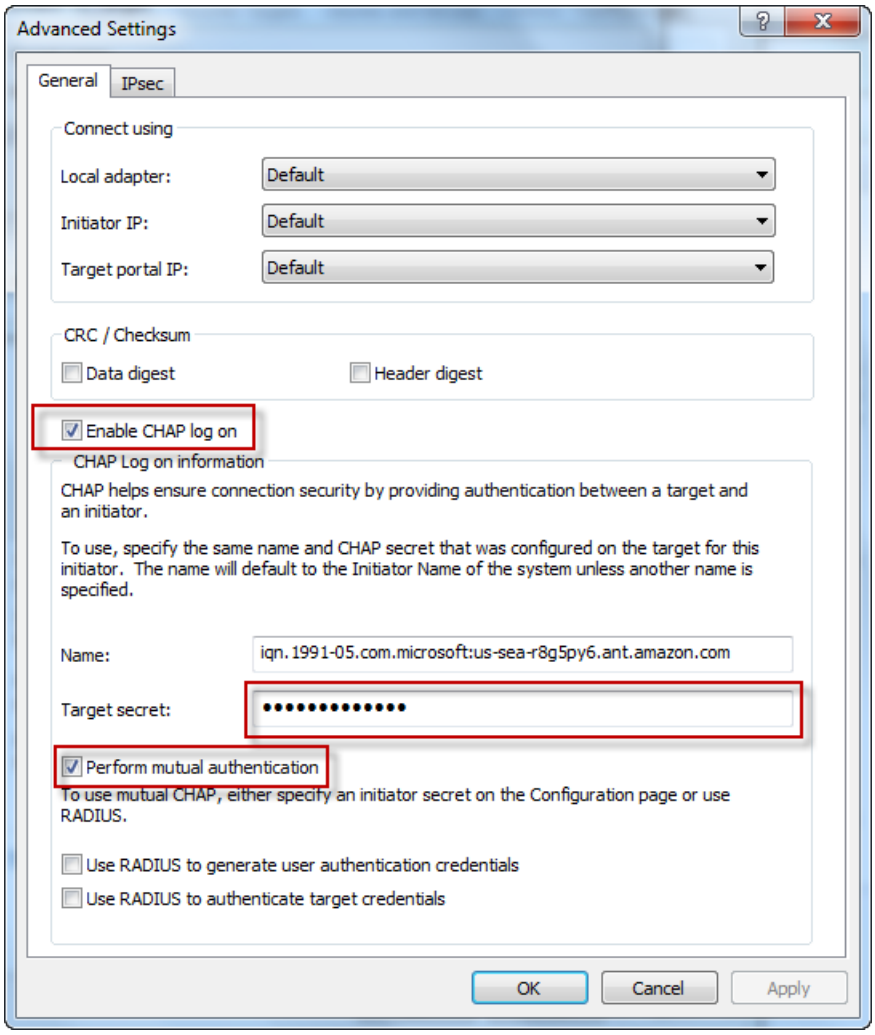
Advanced Settings
General
IPsec
Connect using
Local adapter:
Default
Initiator IP:
Default
Target portal IP:
Default
CRC / Checksum
Data digest
Header digest
Enable CHAP log on
CHAP Log on information
CHAP helps ensure connection security by providing authentication between a target and an initiator.
To use, specify the same name and CHAP secret that was configured on the target for this initiator. The name will default to the Initiator Name of the system unless another name is specified.
Name:
iqn.1991-05.com.microsoft:us-sea-r8g5py6.ant.amazon.com
Target secret:
Perform mutual authentication
To use mutual CHAP, either specify an initiator secret on the Configuration page or use RADIUS.
Use RADIUS to generate user authentication credentials
Use RADIUS to authenticate target credentials
OK
Cancel
Apply

i. Click to select the **Enable CHAP log on** check box.

ii. Type the secret that is required to authenticate the initiator. This secret maps to the secret typed into the **Secret used to Authenticate Initiator** box in the **Configure CHAP Authentication** dialog box. For more information, see Configuring CHAP Authentication for Your Storage Volume (p. 116).

iii. Click to select the **Perform mutual authentication** check box.

iv. Click **OK** to apply the changes.

f. In the **Connect to Target** dialog box, click **OK**.

4. If you provided the correct secret key, the target shows a status of **Connected**.

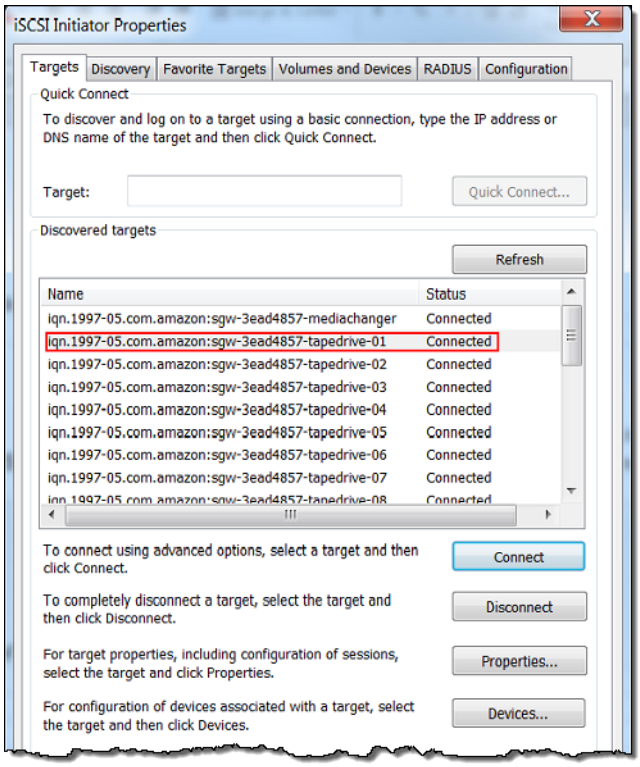


### To configure mutual CHAP on a Red Hat Linux client

In this procedure, you configure CHAP in the Linux iSCSI initiator using the same keys that you used to configure CHAP for the storage volume on the AWS Storage Gateway console.

1. Ensure that the iSCSI daemon is running and that you have already connected to a target. If you have not completed these two tasks, see Connecting from a Red Hat Client to Your Storage Volume (p. 114).
2. Disconnect and remove any existing configuration for the target for which you are about to configure CHAP.

   a. To find the target name and ensure it is a defined configuration, list the saved configurations using the following command.

   ```
   sudo /sbin/iscsiadm --mode node
   ```

   b. Disconnect from the target.

   The following command disconnects from the target named *myvolume* that is defined in the Amazon iSCSI qualified name (IQN). Change the target name and IQN as required for your situation.

   ```
   sudo /sbin/iscsiadm --mode node --logout GATEWAY_IP:3260,1 iqn.1997-05.com.amazon:myvolume
   ```

   c. Remove the configuration for the target.

   The following command removes the configuration for the *myvolume* target.

```
sudo /sbin/iscsiadm --mode node --op delete --targetname iqn.1997-
05.com.amazon:myvolume
```

3. Edit the iSCSI configuration file to enable CHAP.

   a. Get the name of the initiator (that is, the client you are using).

      The following command gets the initiator name from the `/etc/iscsi/initiatorname.iscsi` file.

      ```
      sudo cat /etc/iscsi/initiatorname.iscsi
      ```

      The output from this command will look like this:

      ```
      InitiatorName=iqn.1994-05.com.redhat:8e89b27b5b8
      ```

   b. Open the `/etc/iscsi/iscsid.conf` file.

   c. Uncomment the following lines in the file and specify the correct values for *`username`*, *`password`*, *`username_in`*, and *`password_in`*.

      ```
      node.session.auth.authmethod = CHAP
      node.session.auth.username = username
      node.session.auth.password = password
      node.session.auth.username_in = username_in
      node.session.auth.password_in = password_in
      ```

      To specify the correct values, use the following table as guidance.

| Configuration Setting | Value |
|---|---|
| *`username`* | The initiator name that you found in a previous step in this procedure. The value will start with *iqn*. For example, **`iqn.1994-05.com.redhat:8e89b27b5b8`** is a valid *`username`*. |
| *`password`* | The secret key used to authenticate the initiator (the client you are using) when it communicates with the storage volume. |
| *`username_in`* | The IQN of the target storage volume. The value will start with *iqn* and end with the target name. For example, **`iqn.1997-05.com.amazon:myvolume`** is a valid *`username_in`*. |
| *`password_in`* | The secret key used to authenticate the target (the storage volume) when it communicates to the initiator. |

   d. Save the changes in the configuration file, and then close the file.

4. Discover and log in to the target. To do so, you can follow the steps provided in Connecting from a Red Hat Client to Your Storage Volume (p. 114).

# Performing Volume Gateway Maintenance Tasks

AWS Storage Gateway supports several maintenance tasks after you activate your gateway. This section groups these tasks as follows:

- **Tasks you perform in the AWS Storage Gateway console –** For more information, see Common Gateway Maintenance Tasks (p. 125).
- **Tasks you perform on-premises –** If your gateway is deployed on-premises, you perform most of the maintenance tasks in the local console of your gateway. For more information, see Maintenance Tasks When Running a Gateway On-Premises (p. 137).
- **Tasks you perform when the gateway is running as an Amazon EC2 instance –** For more information, see Maintenance Tasks When Running a Gateway on an Amazon EC2 Instance (p. 160).

**Topics**


- Common Gateway Maintenance Tasks (p. 125)
- Maintenance Tasks When Running a Gateway On-Premises (p. 137)
- Maintenance Tasks When Running a Gateway on an Amazon EC2 Instance (p. 160)
- Managing the Upload Buffer and Cache Storage (Gateway-Cached) (p. 169)
- Managing the Upload Buffer (Gateway-Stored) (p. 178)


## Common Gateway Maintenance Tasks

This section describes gateway-cached and gateway-stored maintenance tasks that you can perform on the AWS Storage Gateway console. For maintenance tasks specific to running a gateway on an Amazon EC2 instance, see Maintenance Tasks When Running a Gateway on an Amazon EC2 Instance (p. 160).

**Topics**


- Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console (p. 127)
- Managing Gateway Updates Using the AWS Storage Gateway Console (p. 129)
- Updating Bandwidth Rate Limits for Your Gateway (p. 130)
- Deleting a Gateway Using the AWS Storage Gateway Console (p. 136)


You can perform many tasks related to gateway maintenance on the **Gateway** tab of the AWS Storage Gateway console. The following screenshot shows the **Gateway** tab.

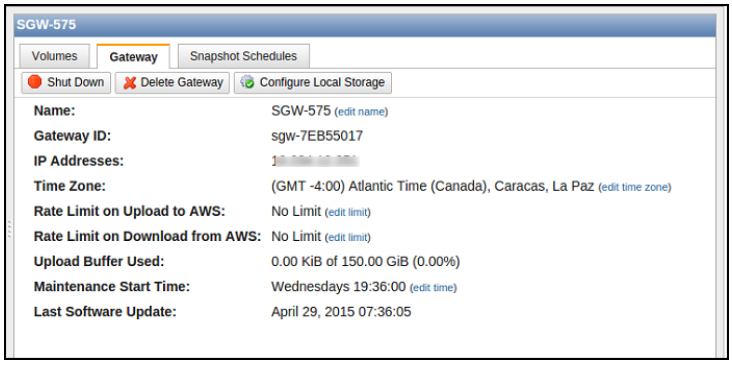


The following table summarizes the information on the **Gateway** tab that you can update. To edit the value for a particular item, click the **edit** link at the end of the box for that item.

| Maintenance Item | Description |
|---|---|
| **Name** | The gateway name. You can change this name as needed. If you use Amazon CloudWatch to view your gateway metrics (see Using Amazon CloudWatch Metrics (p. 186)), you might want to take note of the previous name and the new name to avoid confusion. You can also locate your gateway using the gateway ID, which remains the same through the life of the gateway. |
| **Gateway ID** | The gateway identifier. AWS Storage Gateway assigns a unique identifier for each gateway. This value cannot be changed. |
| **IP Addresses** | The Internet Protocol (IP) addresses using which your storage applications access the gateway's storage volumes. Applications can access volumes using more than one IP address, if the gateway is hosted on a server with more than one network interface card. In this scenario, all addresses that can be used to communicate with the gateway are listed. |
| **Time Zone** | The time zone. AWS Storage Gateway uses this item when displaying time-based information, such as maintenance messages from AWS and snapshot scheduling. |
| **Rate Limit on Upload to AWS** | The limit set for upload throughput from the gateway to AWS. Using this item, you can apply bandwidth throttling to your gateway to control the amount of network bandwidth used for uploads. Specify the rate limit as kilobits per second (kbps). The default is no rate limit on upload. For more information on updating this bandwidth, see Updating Bandwidth Rate Limits for Your Gateway (p. 130). |
| **Rate Limit on Download from AWS** | The limit set for download throughput from AWS to your gateway. Using this item, you can apply bandwidth throttling to your gateway to control the amount of network bandwidth used for downloads. Specify the rate limit in kilobits per second (kbps). The default is no rate limit on download. For more information about updating this bandwidth, see Updating Bandwidth Rate Limits for Your Gateway (p. 130). |
| **Upload Buffer Used** | The upload buffer used. This value cannot be changed. For information about how to monitor the upload buffer and how it changes over time, see Monitoring the Upload Buffer (p. 191). |
| **Maintenance Start Time** | The beginning of the maintenance time period. Each gateway has a maintenance window of one time period per week. During activation, a default time period is assigned to your gateway. To change that maintenance time, click **edit** for this item, and then specify a day of the week and time of the day in the time zone of the gateway. |
| **Last Software Update** | The date and time the last successful software update was applied to the gateway. |
| **Apply Update Now** | A message button. If there is an update for your AWS Storage Gateway installation, a message box displaying this button appears on the console. To apply the update immediately, click **Apply Update Now**. If you don't apply the update immediately, AWS Storage Gateway applies the update based on your **Maintenance Start Time** setting. For more information, see Managing Gateway Updates Using the AWS Storage Gateway Console (p. 129). |

## Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console

This section discusses shutting down and starting a gateway. For example, you might need to shut down your gateway to apply a patch to your hypervisor host. While a gateway is shut down, your applications cannot access storage volumes and therefore cannot write any data to these storage volumes. The gateway also stops uploading any data to AWS when shut down.

Before shutting down your gateway, you must stop any applications that are writing to storage volumes by stopping your iSCSI initiator connection. If a snapshot is in progress when the gateway is shut down, the snapshot will resume on gateway restart. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the AWS Storage Gateway console. For more information, see Editing a Snapshot Schedule (p. 210).

> **Note**
> When you shut down a gateway using the AWS Storage Gateway console, you are stopping the gateway. However, the gateway VM remains on. If you need to turn off the VM, use your VMware vSphere client or Microsoft Hyper-V Manager to connect to your host, and then turn off the specific VM. In most common scenarios in which you use the gateway after activation, you don't need to shut down the gateway VM.

### To shut down a gateway

1. In the **Navigation** pane of the **AWS Storage Gateway** console, select the gateway.
2. Click **Shut Down**.

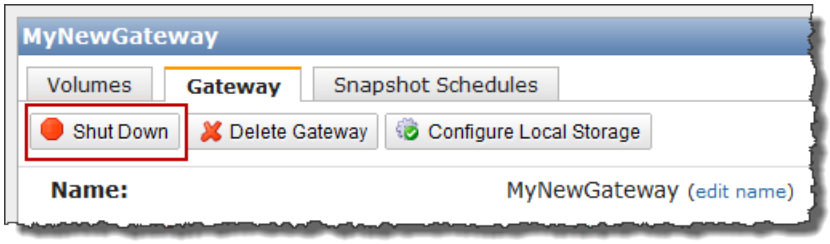


3. In the confirmation dialog box, click **Okay**.

4. In the **Gateway Shut Down** dialog box, click **Close**.

5. While the gateway is shutting down, you might see a message that your gateway is in the process of shutting down.

   The options you have at this point depend on the type of setup you have, gateway-cached or gateway-stored.

- For a gateway-stored setup, you have the option of deleting the gateway. Do not delete the gateway if you plan to restart the gateway and continue working with it.

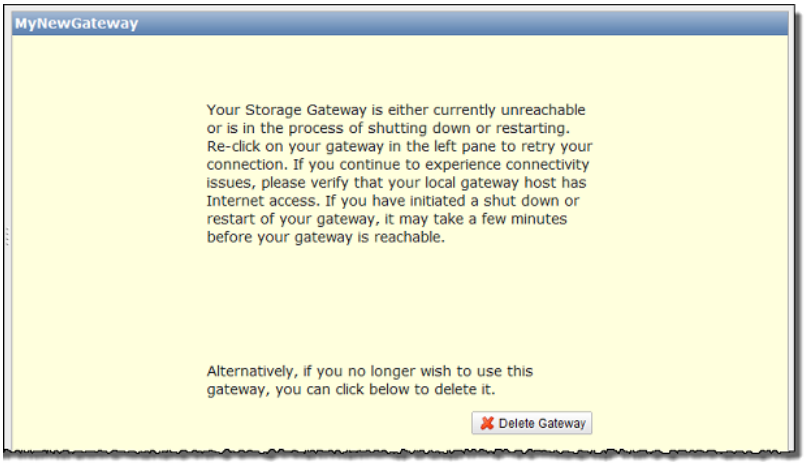
MyNewGateway
Your Storage Gateway is either currently unreachable or is in the process of shutting down or restarting. Re-click on your gateway in the left pane to retry your connection. If you continue to experience connectivity issues, please verify that your local gateway host has Internet access. If you have initiated a shut down or restart of your gateway, it may take a few minutes before your gateway is reachable.
Alternatively, if you no longer wish to use this gateway, you can click below to delete it.
Delete Gateway

- For a gateway-cached setup, you have the option of deleting the gateway or creating a snapshot from a recovery point. For more information on volume recovery points, see Using Recovery Snapshots (Gateway-Cached) (p. 237).

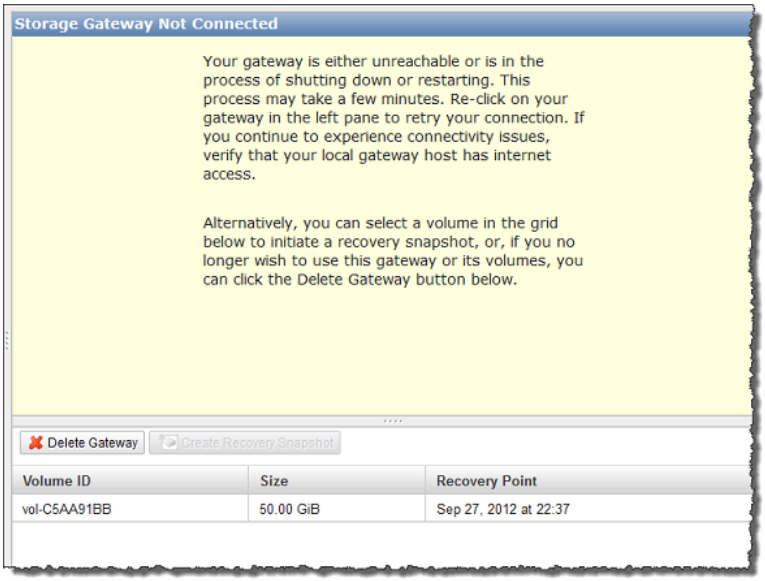
Storage Gateway Not Connected
Your gateway is either unreachable or is in the process of shutting down or restarting. This process may take a few minutes. Re-click on your gateway in the left pane to retry your connection. If you continue to experience connectivity issues, verify that your local gateway host has internet access.
Alternatively, you can select a volume in the grid below to initiate a recovery snapshot, or, if you no longer wish to use this gateway or its volumes, you can click the Delete Gateway button below.
Delete Gateway
Create Recovery Snapshot
Volume ID
Size
Recovery Point
vol-C5AA91BB
50.00 GiB
Sep 27, 2012 at 22:37

6. In the left navigation pane, select the gateway.

   A **Restart** button is displayed.

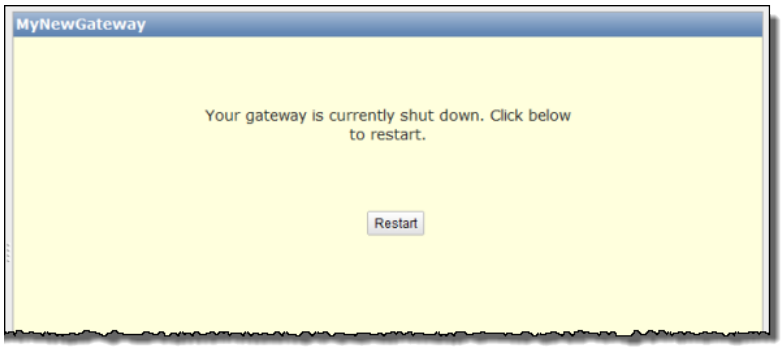


### To turn on a gateway

1. In the **Navigation** pane of the AWS Storage Gateway console, select the gateway to restart.
2. Click **Restart**.

## Managing Gateway Updates Using the AWS Storage Gateway Console

AWS Storage Gateway periodically deploys important updates and patches to your gateway that must be applied. AWS will notify you in advance of any updates to your gateway by using the AWS Storage Gateway console and email. Software updates force a restart of your gateway, which typically takes a few minutes to complete. You do not have to take any action during this update.

During this period, you should not restart the VM manually. After the update, your gateway and its volumes will be in the same states as they were before the update. While the software update is being applied, read and write operations from applications running on the initiators to gateway-storage volume targets are buffered at the iSCSI layer on the initiator. When the update is complete, the gateway processes these operations.

You can choose to let AWS Storage Gateway apply updates according to the maintenance schedule for your gateway, or you can apply the updates yourself. When you deploy and activate your gateway, a default weekly maintenance schedule is set. You can modify this schedule at any time by clicking the **Gateway** tab of the AWS Storage Gateway console and clicking **edit time** next to **Maintenance Start Time**.

The following screenshot shows the **Gateway** tab with a maintenance message and the **Apply Update Now** button. To apply an update yourself, click **Apply Update Now** to start the process. The date and time the last successful software update was applied to the gateway is displayed in the **Gateway** tab.

> **Important**
> A software update forces your gateway to restart. You can minimize the chance of disruption to your applications by increasing your iSCSI initiators' timeouts. For more information about increasing iSCSI initiator timeouts for Windows and Linux, see Customizing Your Windows iSCSI Settings (p. 112) and Customizing Your Linux iSCSI Settings (p. 115).

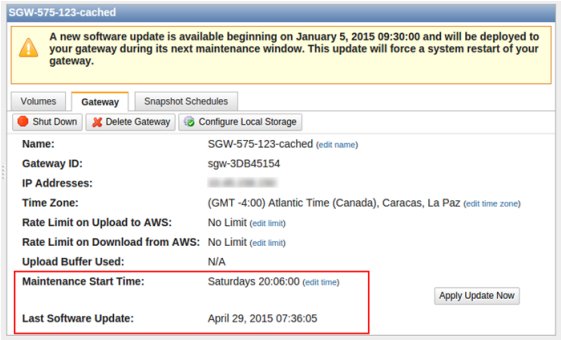
SGW-575-123-cached
A new software update is available beginning on January 5, 2015 09:30:00 and will be deployed to your gateway during its next maintenance window. This update will force a system restart of your gateway.
Volumes
Gateway
Snapshot Schedules
Shut Down
Delete Gateway
Configure Local Storage
Name:
SGW-575-123-cached (edit name)
Gateway ID:
sgw-3DB45154
IP Addresses:
Time Zone:
(GMT -4:00) Atlantic Time (Canada), Caracas, La Paz (edit time zone)
Rate Limit on Upload to AWS:
No Limit (edit limit)
Rate Limit on Download from AWS:
No Limit (edit limit)
Upload Buffer Used:
N/A
Maintenance Start Time:
Saturdays 20:06:00 (edit time)
Apply Update Now
Last Software Update:
April 29, 2015 07:36:05

## Updating Bandwidth Rate Limits for Your Gateway

You can limit (or throttle) the upload throughput from the gateway to AWS or the download throughput from your AWS to your gateway. Using bandwidth throttling helps you to control the amount of network bandwidth used by your gateway. By default, a newly activated gateway has no rate limits on upload or download.

You can specify a rate limit by using the AWS Storage Gateway console, or programmatically using either the AWS Storage Gateway API (see UpdateBandwidthRateLimit) or an AWS Software Development Kit (SDK). The ability to change throttling programmatically lets you change limits automatically throughout the day—for example, by scheduling tasks to change the bandwidth. This section describes how you can change these limits by using the AWS Storage Gateway console.

For examples of changing bandwidth rate limits programmatically, see the following topics: Updating Gateway Bandwidth Rate Limits Using the AWS SDK for Java (p. 131), Updating Gateway Bandwidth Rate Limits Using the AWS SDK for .NET (p. 133), and Updating Gateway Bandwidth Rate Limits Using the AWS Tools for Windows PowerShell (p. 135).

**Topics**

- Updating Gateway Bandwidth Rate Limits Using the AWS SDK for Java (p. 131)
- Updating Gateway Bandwidth Rate Limits Using the AWS SDK for .NET (p. 133)
- Updating Gateway Bandwidth Rate Limits Using the AWS Tools for Windows PowerShell (p. 135)

### To change a gateway's bandwidth throttling using the console

1. In the **Navigation** pane of the **AWS Storage Gateway** console, select the gateway you want to manage.
2. In the right pane, click the **Gateway** tab.
3. Next to the limit you want to change, click **edit limit**.

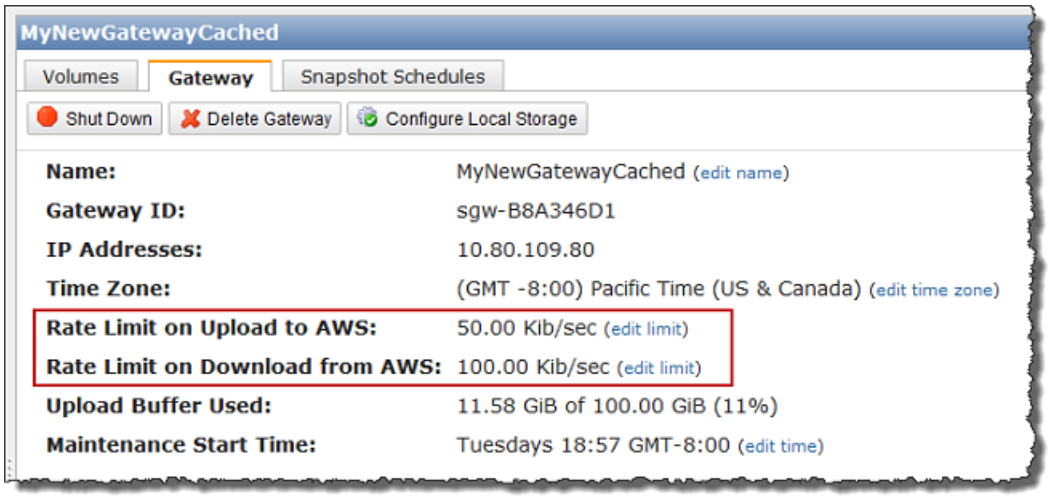


4. In the **Edit Rate Limits** dialog box, enter new limit values, and then click **Save**.

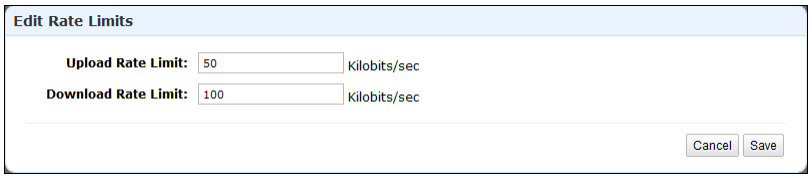


## Updating Gateway Bandwidth Rate Limits Using the AWS SDK for Java

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth rate limits using the AWS SDK for Java. To use the example code, you should be familiar with running a Java console application. For more information, go to Getting Started in the *AWS SDK for Java Developer Guide*.

### Example : Updating Gateway Bandwidth Limits Using the AWS SDK for Java

The following Java code example updates a gateway's bandwidth rate limits. You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints you can use with AWS Storage Gateway, go to Regions and Endpoints in the *AWS General Reference*.

```
import java.io.IOException;

import com.amazonaws.AmazonClientException;
import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.storagegateway.AWSStorageGatewayClient;
import com.amazonaws.services.storagegateway.model.UpdateBandwidthRateLimitRe
quest;
import com.amazonaws.services.storagegateway.model.UpdateBandwidthRateLimitRes
ult;


public class UpdateBandwidthExample {

    public static AWSStorageGatewayClient sgClient;

    // The gatewayARN
    public static String gatewayARN = "*** provide gateway ARN ***";

    // The endpoint
   static String serviceURL = "https://storagegateway.us-east-1.amazonaws.com";


    // Rates
    static long uploadRate = 51200;  // Bits per second, minimum 51200
    static long downloadRate = 102400;   // Bits per second, minimum 102400

    public static void main(String[] args) throws IOException {

        // Create a storage gateway client
        sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(
                ListDeleteVolumeSnapshotsExample.class.getResourceAs
Stream("AwsCredentials.properties")));
        sgClient.setEndpoint(serviceURL);

        UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

    }

    private static void UpdateBandwidth(String gatewayARN2, long uploadRate2,
            long downloadRate2) {
        try
        {
            UpdateBandwidthRateLimitRequest updateBandwidthRateLimitRequest =
                new UpdateBandwidthRateLimitRequest()
                .withGatewayARN(gatewayARN)
                .withAverageDownloadRateLimitInBitsPerSec(downloadRate)
                .withAverageUploadRateLimitInBitsPerSec(uploadRate);

            UpdateBandwidthRateLimitResult updateBandwidthRateLimitResult =
sgClient.updateBandwidthRateLimit(updateBandwidthRateLimitRequest);
           String returnGatewayARN = updateBandwidthRateLimitResult.getGatewa
```

```
yARN();
            System.out.println("Updated the bandwidth rate limits of " + re
turnGatewayARN);
            System.out.println("Upload bandwidth limit = " + uploadRate + "
bits per second");
            System.out.println("Download bandwidth limit = " + downloadRate +
" bits per second");
        }
        catch (AmazonClientException ex)
        {
            System.err.println("Error updating gateway bandwith.\n" + ex.to
String());
        }
    }
}
```

### Updating Gateway Bandwidth Rate Limits Using the AWS SDK for .NET

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth rate limits using the AWS Software Development Kit (SDK) for .NET. To use the example code, you should be familiar with running a .NET console application. For more information, go to Getting Started in the *AWS SDK for .NET Developer Guide*.

### Example : Updating Gateway Bandwidth Limits Using the AWS SDK for .NET

The following C# code example updates a gateway's bandwidth rate limits. You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints you can use with AWS Storage Gateway, go to Regions and Endpoints in the *AWS General Reference*.

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using Amazon.StorageGateway;
using Amazon.StorageGateway.Model;

namespace AWSStorageGateway
{
    class UpdateBandwidthExample
    {
        static AmazonStorageGatewayClient sgClient;
        static AmazonStorageGatewayConfig sgConfig;

        // The gatewayARN
        public static String gatewayARN = "*** provide gateway ARN ***";

        // The endpoint
        static String serviceURL = "https://storagegateway.us-east-1.amazon
aws.com";

        // Rates
        static long uploadRate = 51200;  // Bits per second, minimum 51200
       static long downloadRate = 102400;   // Bits per second, minimum 102400


        public static void Main(string[] args)
        {
            // Create a storage gateway client
            sgConfig = new AmazonStorageGatewayConfig();
            sgConfig.ServiceURL = serviceURL;
            sgClient = new AmazonStorageGatewayClient(sgConfig);

            UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

            Console.WriteLine("\nTo continue, press Enter.");
            Console.Read();
        }

        public static void UpdateBandwidth(string gatewayARN, long uploadRate,
 long downloadRate)
        {
            try
            {
               UpdateBandwidthRateLimitRequest updateBandwidthRateLimitRequest
 =
                    new UpdateBandwidthRateLimitRequest()
                    .WithGatewayARN(gatewayARN)
                    .WithAverageDownloadRateLimitInBitsPerSec(downloadRate)
                    .WithAverageUploadRateLimitInBitsPerSec(uploadRate);
```

```
             UpdateBandwidthRateLimitResponse updateBandwidthRateLimitResponse
 = sgClient.UpdateBandwidthRateLimit(updateBandwidthRateLimitRequest);
                String returnGatewayARN = updateBandwidthRateLimitResponse.Up
dateBandwidthRateLimitResult.GatewayARN;
                Console.WriteLine("Updated the bandwidth rate limits of " +
returnGatewayARN);
                Console.WriteLine("Upload bandwidth limit = " + uploadRate + "
 bits per second");
                Console.WriteLine("Download bandwidth limit = " + downloadRate
 + " bits per second");
            }
            catch (AmazonStorageGatewayException ex)
            {
                Console.WriteLine("Error updating gateway bandwith.\n" +
ex.ToString());
            }
        }
    }
}
```

## Updating Gateway Bandwidth Rate Limits Using the AWS Tools for Windows PowerShell

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth rate limits using the AWS Tools for Windows PowerShell. To use the example code, you should be familiar with running a PowerShell script. For more information, go to Getting Started in the AWS Tools for Windows PowerShell User Guide.

### Example : Updating Gateway Bandwidth Limits Using the AWS Tools for Windows PowerShell

The following PowerShell script example updates a gateway's bandwidth rate limits. You need to update the script and provide your gateway Amazon Resource Name (ARN) and the upload and download limits.

```
<#
.DESCRIPTION
    Update Gateway bandwidth limits.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from http://console.amazonaws.cn/powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
   For more info see, http://docs.amazonaws.cn/powershell/latest/userguide/spe
cifying-your-aws-credentials.html

.EXAMPLE
    powershell.exe .\SG_UpdateBandwidth.ps1
#>

$UploadBandwidthRate = 51200
$DownloadBandwidthRate = 102400
$gatewayARN = "*** provide gateway ARN ***"

#Update Bandwidth Rate Limits
Update-SGBandwidthRateLimit -GatewayARN $gatewayARN `
                            -AverageUploadRateLimitInBitsPerSec $UploadBandwid
thRate `
                           -AverageDownloadRateLimitInBitsPerSec $DownloadBand
widthRate

$limits =  Get-SGBandwidthRateLimit -GatewayARN $gatewayARN

Write-Output("`nGateway: " + $gatewayARN);
Write-Output("`nNew Upload Rate: " + $limits.AverageUploadRateLimitInBitsPerSec)
Write-Output("`nNew Download Rate: " + $limits.AverageDownloadRateLimitInBitsPer
Sec)
```

## Deleting a Gateway Using the AWS Storage Gateway Console

Deleting a gateway removes the gateway as an activated gateway that you can use to store application data. The deleted gateway no longer shows on the AWS Storage Gateway console, and any existing iSCSI connections you have open to the gateway will be closed. Reusing the VM for a new gateway is not supported.

> **Important**
> You no longer pay software charges after the gateway is deleted. However, your existing Amazon EBS snapshots persist, and you will continue to be billed for these snapshots. You can choose to remove all remaining Amazon EBS snapshots by canceling your Amazon EC2 subscription. If you prefer not to cancel your Amazon EC2 subscription, you can delete your snapshots using the Amazon EC2 console. For more information, go to the AWS Storage Gateway details page.

### To delete a gateway

1. In the **Navigation** pane of the **AWS Storage Gateway** console, select the gateway you want to delete.
2. On the **Gateway** tab, click **Delete Gateway**.

**Important**
Be sure that there are no applications currently writing to the gateway's volumes. If you delete the gateway while it is in use, data loss can occur.

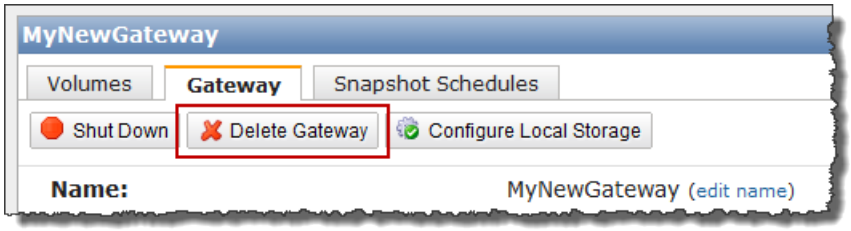


3. To confirm the deletion, click **OK**.

   At this point, the deleted gateway is no longer an activated gateway. However, the gateway VM still exists in your virtualization environment. To remove the VM, use the VMware vSphere client or Microsoft Hyper-V Manager to connect to the host and remove the VM.

# Maintenance Tasks When Running a Gateway On-Premises

For a gateway deployed on-premises, this section describes the maintenance tasks you can perform using the local gateway console. For information about maintenance tasks when running a gateway on an Amazon EC2 instance, see Maintenance Tasks When Running a Gateway on an Amazon EC2 Instance (p. 160).

**Topics**


- Logging In to Your AWS Storage Gateway Local Console (p. 137)
- Routing Your On-Premises Gateway Through a Proxy (p. 141)
- Configuring Your Gateway to Use a Static IP Address (p. 142)
- Testing Your Gateway Connection to the Internet (p. 145)
- Synchronizing Your Gateway VM Time (p. 146)
- Running Storage Gateway Commands on the Local Console (p. 148)
- Updating Your System (p. 149)
- Viewing Your Gateway System Resource Errors (p. 150)
- Configuring Your Gateway for Multiple NICs (p. 152)
- Creating a Storage Volume on a Gateway with Multiple Network Adapters (p. 157)
- Optimizing Gateway Performance (p. 158)


## Logging In to Your AWS Storage Gateway Local Console

Some gateway maintenance tasks require that you log in to your gateway's local console. You can access the local console through your hypervisor client software. The user is *sguser* and the password is *sgpassword*. These login credentials give you access to menus where you can configure gateway network settings.

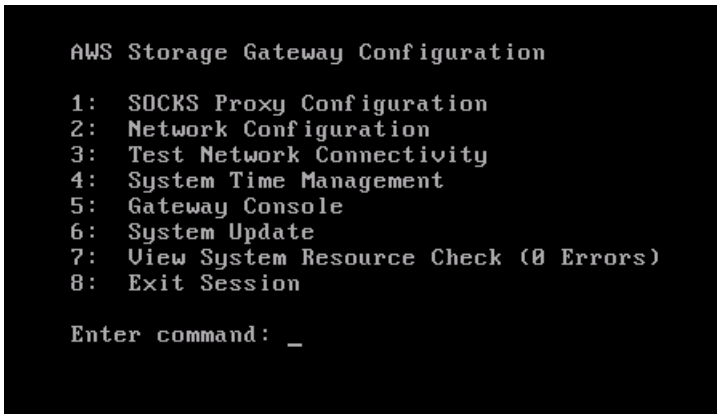


In this topic, we show you how to access the local console of a gateway hosted in these hypervisors:

- VMware ESXi—for more information, see To access your gateway's local console (VMware ESXi) (p. 138).
- Microsoft Hyper-V—for more information, see To access your gateway's local console (Microsoft Hyper-V) (p. 139).

After you access the console, you log in to it. For more information on logging in, see To log in to the gateway local console (p. 140).

### To access your gateway's local console (VMware ESXi)

1. In the VMware vSphere client, select your gateway VM.
2. Ensure that the gateway is turned on.

   **Note**
   If your gateway VM is turned on, a green arrow icon appears with the VM icon, as shown in the following screenshot. If your gateway VM is not turned on, you can turn it on by clicking the green **Power On** icon on the **Toolbar** menu.

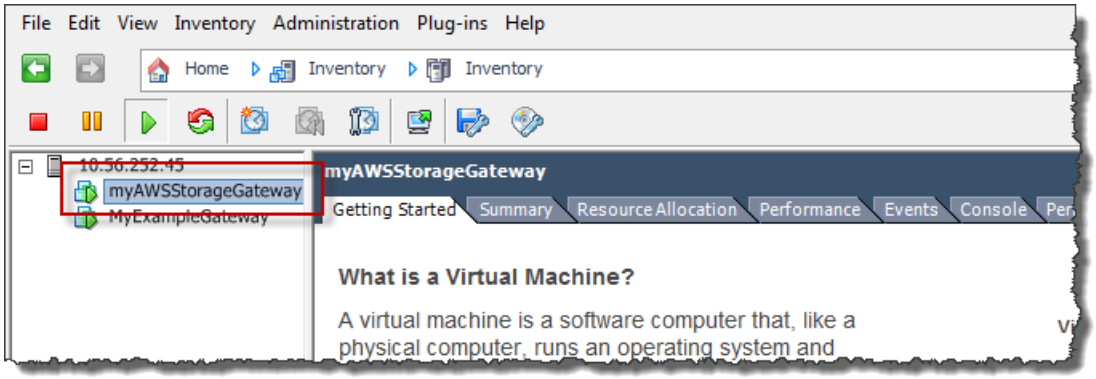


3. Click the **Console** tab.

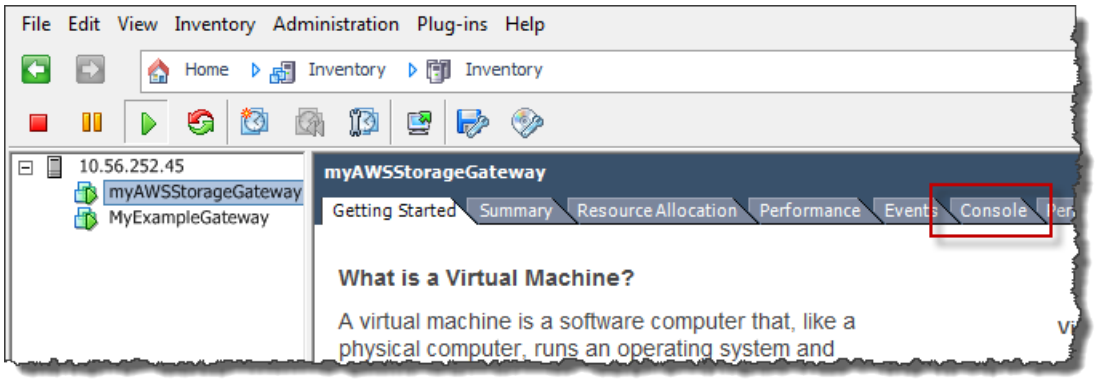
File Edit View Inventory Administration Plug-ins Help
Home Inventory Inventory
10.56.252.45
myAWSStorageGateway
MyExampleGateway
myAWSStorageGateway
Getting Started Summary Resource Allocation Performance Events Console
What is a Virtual Machine?
A virtual machine is a software computer that, like a physical computer, runs an operating system and

4. After a few moments, the VM is ready for you to log in.

   **Note**
   To release the cursor from the console window, press **Ctrl+Alt**.

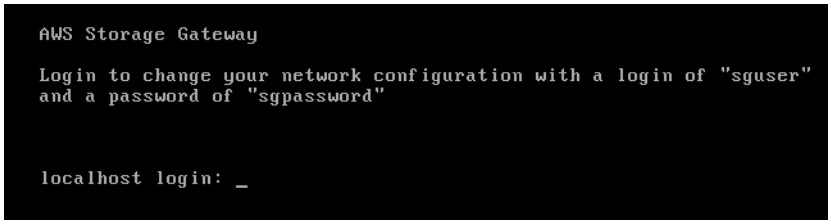
AWS Storage Gateway
Login to change your network configuration with a login of "sguser" and a password of "sgpassword"
localhost login: _

5. To log in, continue to the procedure To log in to the gateway local console (p. 140).

### To access your gateway's local console (Microsoft Hyper-V)

1. In the **Virtual Machines** list of the Microsoft Hyper-V Manager, select your gateway VM.
2. Ensure the gateway is turned on.

   **Note**
   If your gateway VM is turned on, `Running` is displayed as the **State** of the VM, as shown in the following screenshot. If your gateway VM is not turned on, you can turn it on by clicking **Start** in the **Actions** pane.

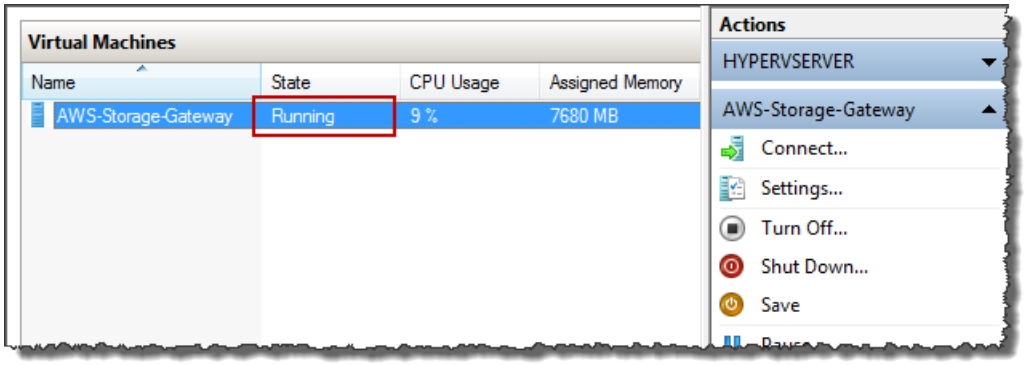


3. In the **Actions** pane, select **Connect**.

   The **Virtual Machine Connection** window appears. If an authentication window appears, type the user name and password provided to you by the hypervisor administrator.

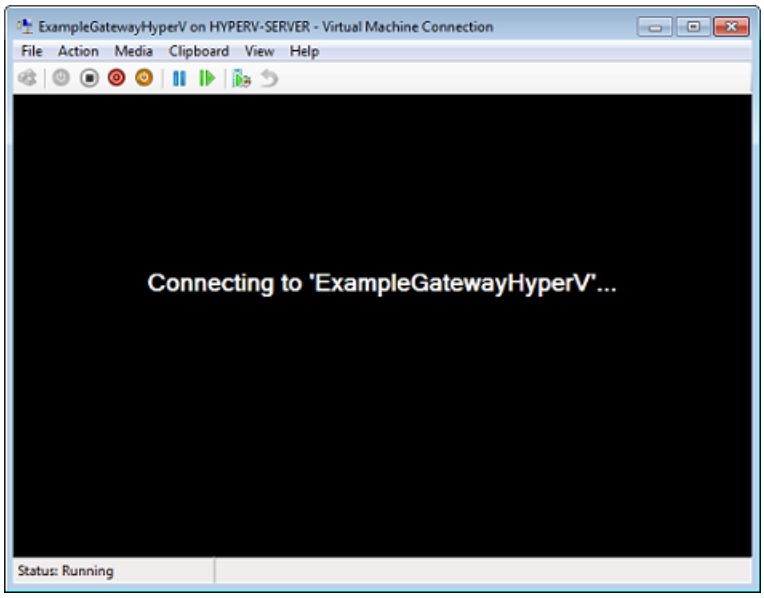
ExampleGatewayHyperV on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
Connecting to 'ExampleGatewayHyperV'...
Status: Running

4. After a few moments, the VM is ready for you to log in.

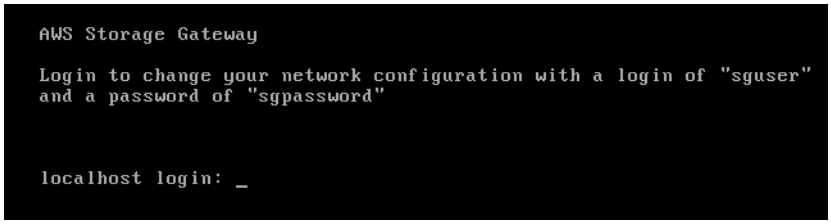
AWS Storage Gateway
Login to change your network configuration with a login of "sguser"
and a password of "sgpassword"
localhost login: _

5. To log in, continue to the procedure To log in to the gateway local console (p. 140).

### To log in to the gateway local console

1. In the login screen, log in to the VM with the user name and password specified in the console window.
2. After you log in, you will see the **AWS Storage Gateway Configuration** main menu.

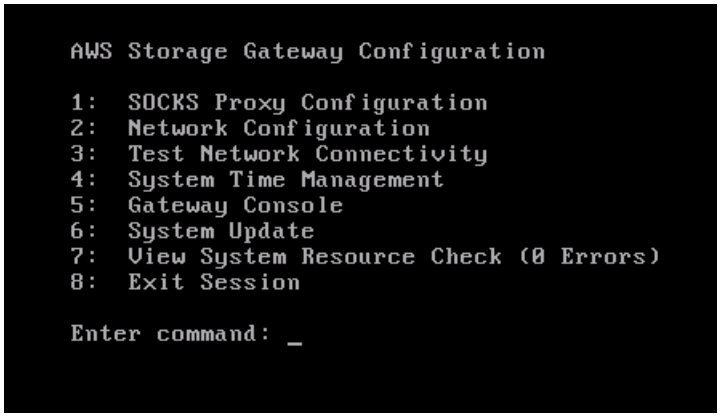

| To... | See... |
|---|---|
| Configure a SOCKS proxy for your gateway | Routing Your On-Premises Gateway Through a Proxy (p. 141). |
| Configure static IP addresses for your gateway's interfaces | Configuring Your Gateway to Use a Static IP Address (p. 142). |
| Test network connectivity | Testing Your Gateway Connection to the Internet (p. 145). |
| Manage VM time | Synchronizing Your Gateway VM Time (p. 146). |
| Run Storage Gateway console commands | Running Storage Gateway Commands on the Local Console (p. 148). |
| Update your system | Updating Your System (p. 149). |

To exit the configuration session, type **7** to exit the menu.

## Routing Your On-Premises Gateway Through a Proxy

AWS Storage Gateway supports the configuration of a Socket Secure version 5 (SOCKS5) proxy between your on-premises gateway and AWS.

**Note**
The only proxy configuration AWS Storage Gateway supports is SOCKS5.

If your gateway must use a proxy server to communicate to the Internet, then you need to configure the SOCKS proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy. After you do so, AWS Storage Gateway will route all HyperText Transfer Protocol Secure (HTTPS) traffic through your proxy server.

### To route your gateway Internet traffic through a local proxy server

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 137).
2. On the **AWS Storage Gateway Configuration** main menu, type **1** to begin configuring the SOCKS proxy.

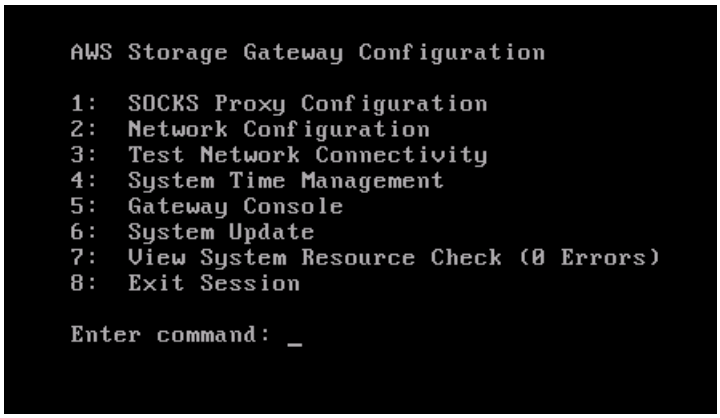


3. Choose one of the following options on the **AWS Storage Gateway SOCKS Proxy Configuration** menu:

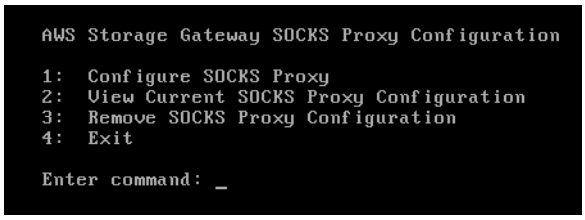
AWS Storage Gateway SOCKS Proxy Configuration
1: Configure SOCKS Proxy
2: View Current SOCKS Proxy Configuration
3: Remove SOCKS Proxy Configuration
4: Exit
Enter command: _

| To... | Do This... |
|---|---|
| Configure a SOCKS proxy | Type option **1**.<br><br>You will need to supply a host name and port to complete configuration. |
| View the current SOCKS proxy configuration | Type option **2**.<br><br>If a SOCKS proxy is not configured, the message `SOCKS Proxy not configured` is displayed. If a SOCKS proxy is configured, the host name and port of the proxy are displayed. |
| Remove a SOCKS proxy configuration | Type option **3**.<br><br>The message `SOCKS Proxy Configuration Removed` is displayed. |
| Exit this menu and return to the previous menu | Type option **4**. |

## Configuring Your Gateway to Use a Static IP Address

The default network configuration for the gateway is Dynamic Host Configuration Protocol (DHCP). With DHCP, your gateway is automatically assigned an IP address. In some cases, you might need to manually assign your gateway's IP as a static IP address. This topic explains how to do so.

### To configure your gateway to use static IP addresses

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 137).
2. On the **AWS Storage Gateway Configuration** main menu, type option **2** to begin configuring a static IP address.

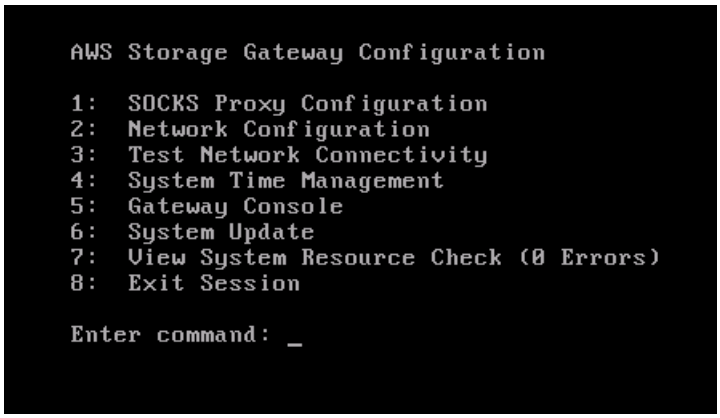


3. Choose one of the following options on the **AWS Storage Gateway Static IP Address Configuration** menu:

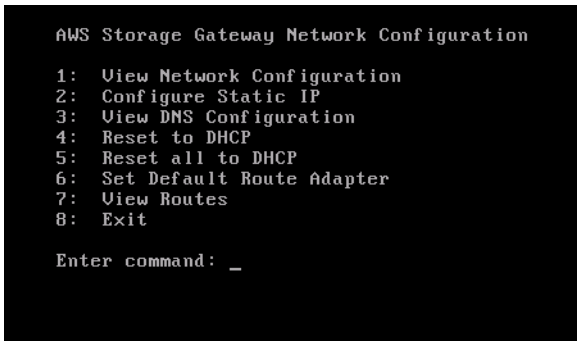


| To... | Do this... |
|---|---|
| View your gateway's network configuration | Type option **1**.<br><br>A list of adapter names appears, and you are prompted to type an adapter name—for example, **`eth0`**. If the adapter you specify is in use, the following information about the adapter is displayed:<br><br>• Media access control (MAC) address<br>• IP address<br>• Netmask<br>• Gateway IP address<br>• DHCP enabled status<br><br>You use the same adapter name when you configure a static IP address (option **2**) as when you set your gateway's default route adapter (option **5**). |

| To... | Do this... |
| --- | --- |
| Configure a static IP address for your gateway | Type option **2**.<br><br>You are prompted to type the following information to configure a static IP:<br><br>• Network adapter name<br>• IP address<br>• Netmask<br>• Default gateway address<br>• Primary Domain Name Service (DNS) address<br>• Secondary DNS address<br><br>**Important**<br>If your gateway has already been activated, you must shut it down and restart it from the AWS Storage Gateway console for the settings to take effect. For more information, see Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console (p. 127).<br><br>If your gateway uses more than one network interface, you must set all enabled interfaces to use DHCP or static IP addresses.<br><br>For example, suppose your gateway VM uses two interfaces configured as DHCP. If you later set one interface to a static IP, the other interface is disabled. To enable the interface in this case, you must set it to a static IP.<br><br>If both interfaces are initially set to use static IP addresses and you then set the gateway to use DHCP, both interfaces will use DHCP. |
| View your gateway's DNS configuration | Type option **3**.<br><br>The IP addresses of the primary and secondary DNS name servers are displayed. |
| Reset your gateway's network configuration to DHCP | Type option **4**.<br><br>All network interfaces are set to use DHCP.<br><br>**Important**<br>If your gateway has already been activated, you must shut down and restart your gateway from the AWS Storage Gateway console for the settings to take effect. For more information, see Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console (p. 127). |

| To... | Do this... |
|---|---|
| Set your gateway's default route adapter | Type option **5**.<br><br>The available adapters for your gateway are shown, and you are prompted to select one of the adapters—for example, **`eth0`**. |
| View routing tables | Type option **6**.<br><br>The default route of your gateway is displayed. |
| Exit this menu and return to the previous menu | Type option **7**. |

## Testing Your Gateway Connection to the Internet

You can use your gateway's local console to test your Internet connection. This test can be useful when you are troubleshooting network issues with your gateway. The network connectivity test does not take into account any SOCKS proxy you might have configured.

### To test your gateway's connection to the Internet

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 137).
2. On the **AWS Storage Gateway Configuration** main menu, type option **3** to begin testing network connectivity.

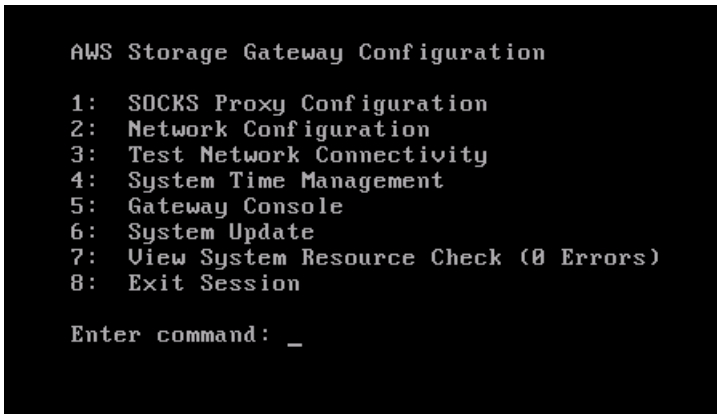


   The console displays the available regions.
3. Select the region you want to test, as shown in the following screenshot.

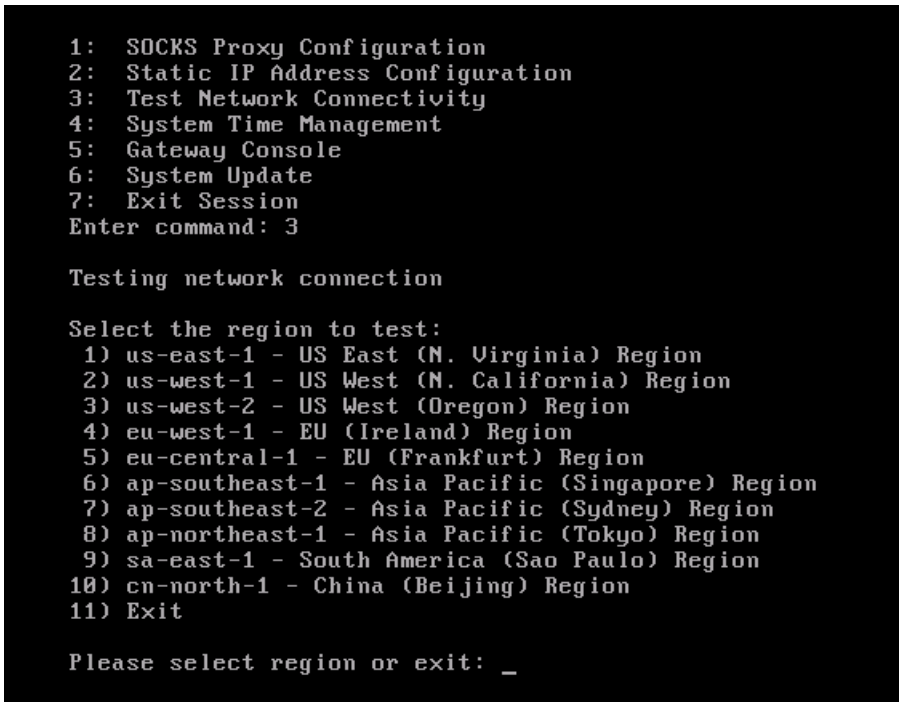
1: SOCKS Proxy Configuration
2: Static IP Address Configuration
3: Test Network Connectivity
4: System Time Management
5: Gateway Console
6: System Update
7: Exit Session
Enter command: 3
Testing network connection
Select the region to test:
1) us-east-1 - US East (N. Virginia) Region
2) us-west-1 - US West (N. California) Region
3) us-west-2 - US West (Oregon) Region
4) eu-west-1 - EU (Ireland) Region
5) eu-central-1 - EU (Frankfurt) Region
6) ap-southeast-1 - Asia Pacific (Singapore) Region
7) ap-southeast-2 - Asia Pacific (Sydney) Region
8) ap-northeast-1 - Asia Pacific (Tokyo) Region
9) sa-east-1 - South America (Sao Paulo) Region
10) cn-north-1 - China (Beijing) Region
11) Exit
Please select region or exit: _

Each endpoint in the selected region displays either a PASSED or FAILED message, as shown following.

| Message | Description |
|---|---|
| [ PASSED ] | AWS Storage Gateway has Internet connectivity. |
| [ FAILED ] | AWS Storage Gateway does not have Internet connectivity. |

## Synchronizing Your Gateway VM Time

After your gateway is deployed and running, in some scenarios the gateway VM's time can drift. For example, if there is a prolonged network outage and your hypervisor host and gateway do not get time updates, then the gateway VM's time will be different from the true time. When there is a time drift, a discrepancy occurs between the stated times when operations such as snapshots occur and the actual times that the operations occur.

For a gateway deployed on VMware ESXi, setting the hypervisor host time and synchronizing the VM time to the host is sufficient to avoid time drift. For more information, see Step 2.2.3: Synchronize VM Time with Host Time (p. 20). For a gateway deployed on Microsoft Hyper-V, you should periodically check your VM's time using the procedure described following.

### To view and synchronize the time of a Hyper-V gateway VM to an NTP server

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 137).
2. On the **AWS Storage Gateway Configuration** main menu, type **4** for **System Time Management**.

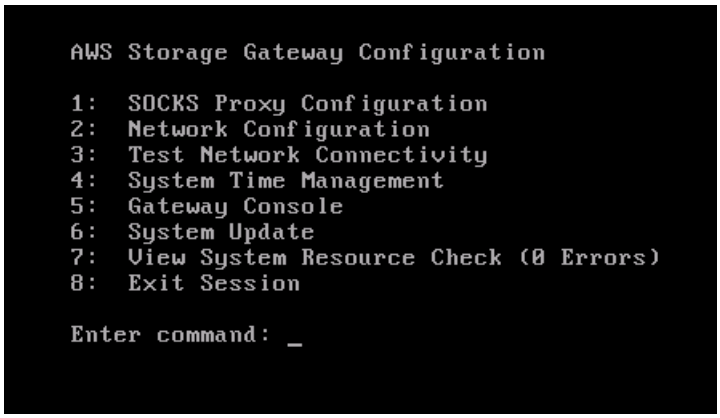


3. On the **System Time Management** menu, type **1** for **View and Synchronize System Time**.

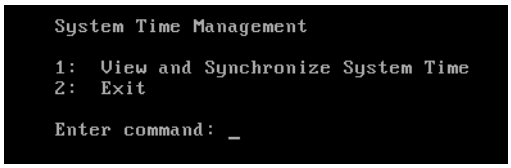


4. If the result indicates that you should synchronize your VM's time to the Network Time Protocol (NTP) time, type **y**. Otherwise, type **n**.

   If you type **y** to synchronize, the synchronization might take a few moments.

   The following screenshot shows a VM that does not require time synchronization.

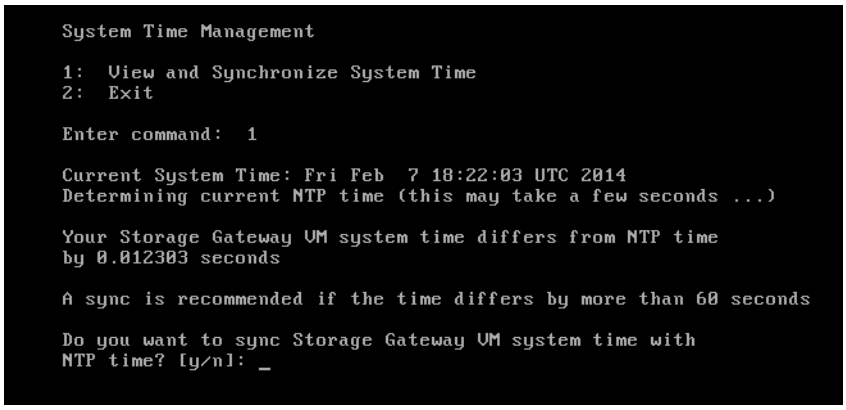
System Time Management
1: View and Synchronize System Time
2: Exit
Enter command: 1
Current System Time: Fri Feb 7 18:22:03 UTC 2014
Determining current NTP time (this may take a few seconds ...)
Your Storage Gateway VM system time differs from NTP time
by 0.012303 seconds
A sync is recommended if the time differs by more than 60 seconds
Do you want to sync Storage Gateway VM system time with
NTP time? [y/n]: _

   The following screenshot shows a VM that does require time synchronization.

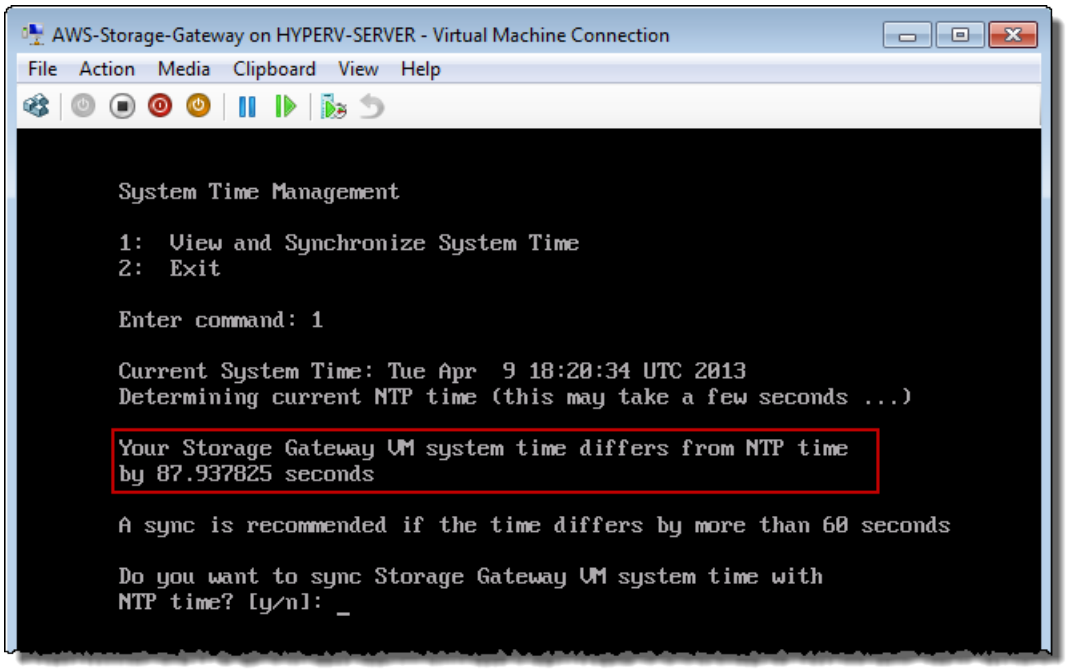
AWS-Storage-Gateway on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
System Time Management
1: View and Synchronize System Time
2: Exit
Enter command: 1
Current System Time: Tue Apr 9 18:20:34 UTC 2013
Determining current NTP time (this may take a few seconds ...)
Your Storage Gateway VM system time differs from NTP time
by 87.937825 seconds
A sync is recommended if the time differs by more than 60 seconds
Do you want to sync Storage Gateway VM system time with
NTP time? [y/n]: _

## Running Storage Gateway Commands on the Local Console

The AWS Storage Gateway console helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the console commands, you can perform maintenance tasks such as saving routing tables or connecting to AWS Support.

### To run a configuration or diagnostic command

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 137).
2. On the **AWS Storage Gateway Configuration** main menu, type option **5** for **Gateway Console**.

3. On the AWS Storage Gateway console, type **h**, and then press the **Return** key.

The console displays the **Available Commands** menu with the available commands and after the menu a **Gateway Console** prompt, as shown in the following screenshot.

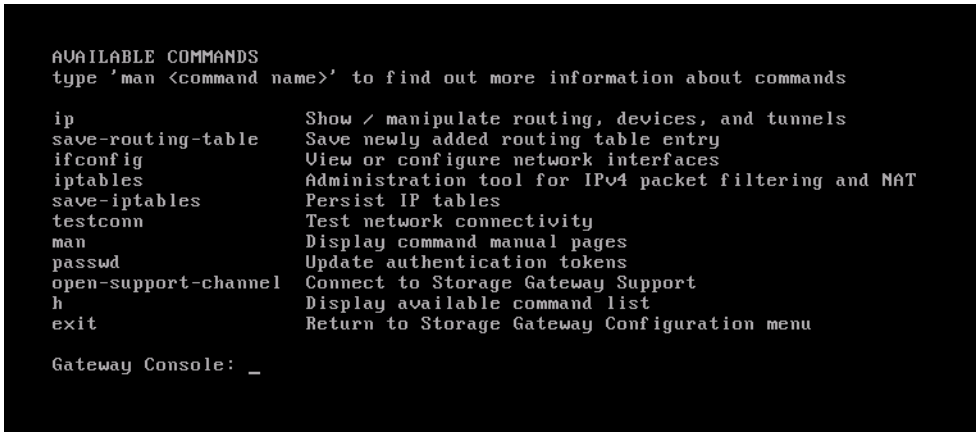
AVAILABLE COMMANDS
type 'man <command name>' to find out more information about commands
ip Show / manipulate routing, devices, and tunnels
save-routing-table Save newly added routing table entry
ifconfig View or configure network interfaces
iptables Administration tool for IPv4 packet filtering and NAT
save-iptables Persist IP tables
testconn Test network connectivity
man Display command manual pages
passwd Update authentication tokens
open-support-channel Connect to Storage Gateway Support
h Display available command list
exit Return to Storage Gateway Configuration menu
Gateway Console: _

4. To learn about a command, type **man** + ***command name*** at the **Gateway Console** prompt.

## Updating Your System

You can use the local console to update the software on your system.

**Warning**
Applying system updates when the gateway is turned on can cause volumes to move to the PASS THROUGH state. Shut down the gateway before applying system updates.

### To update your system software by using the local console

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 137).
2. On the **AWS Storage Gateway Configuration** main menu, type option **6** to begin updating.

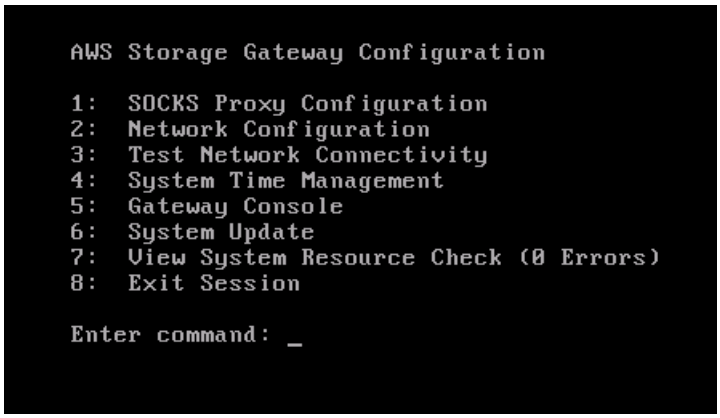


The console displays the **Apply System Updates** menu, as shown in the following screenshot.

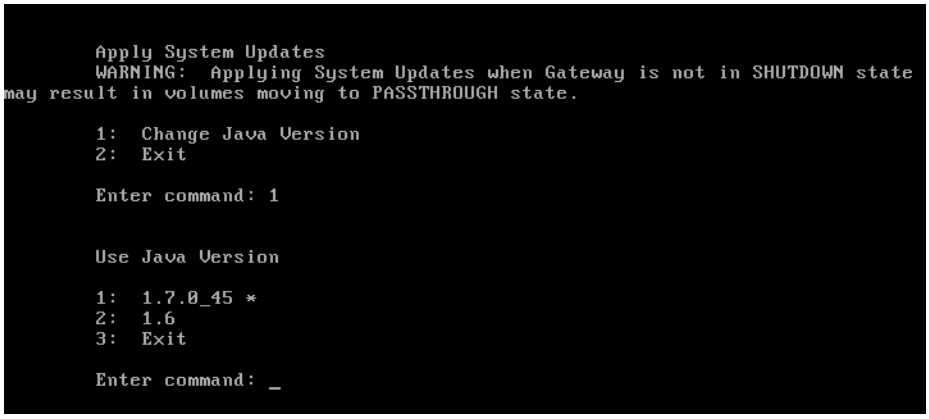
Apply System Updates
WARNING: Applying System Updates when Gateway is not in SHUTDOWN state may result in volumes moving to PASSTHROUGH state.
1: Change Java Version
2: Exit
Enter command: 1
Use Java Version
1: 1.7.0_45 *
2: 1.6
3: Exit
Enter command: _

3. Type one of the following options to use a new version of Java or to exit.

| To... | Do this... |
|---|---|
| Use Java Version 1.7.0_45 * | Type **1**. |
| Use Java Version 1.6 | Type **2**. |
| Exit to the previous menu | Type **3**. |

## Viewing Your Gateway System Resource Errors

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM and determines whether these system resources are sufficient for your gateway to function properly. You can view the results of this check on the gateway's local console.

### To view the status of a system resource check

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 137).
2. In the **AWS Storage Gateway Configuration** main menu, type **7** to view the results of a system resource check.

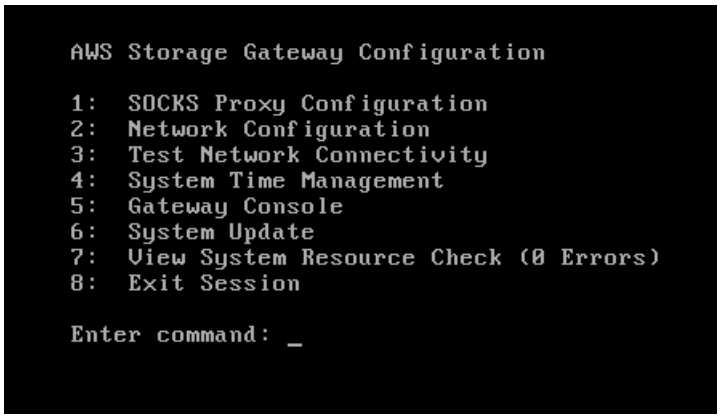


The console displays an **[OK]**, **[WARNING]**, or **[FAIL]** message for each resource as described in the table following.

| Message | Description |
|---|---|
| **[OK]** | The resource has passed the system resource check. |
| **[WARNING]** | The resource does not meet the recommended requirements, but your gateway will continue to function. AWS Storage Gateway displays a message that describes the results of the resource check. |
| **[FAIL]** | The resource does not meet the minimum requirements. Your gateway might not function properly. AWS Storage Gateway displays a message that describes the results of the resource check. |

The console also displays the number of errors and warnings next to the resource check menu option.

The following screenshot shows a **[FAIL]** message and the accompanying error message.

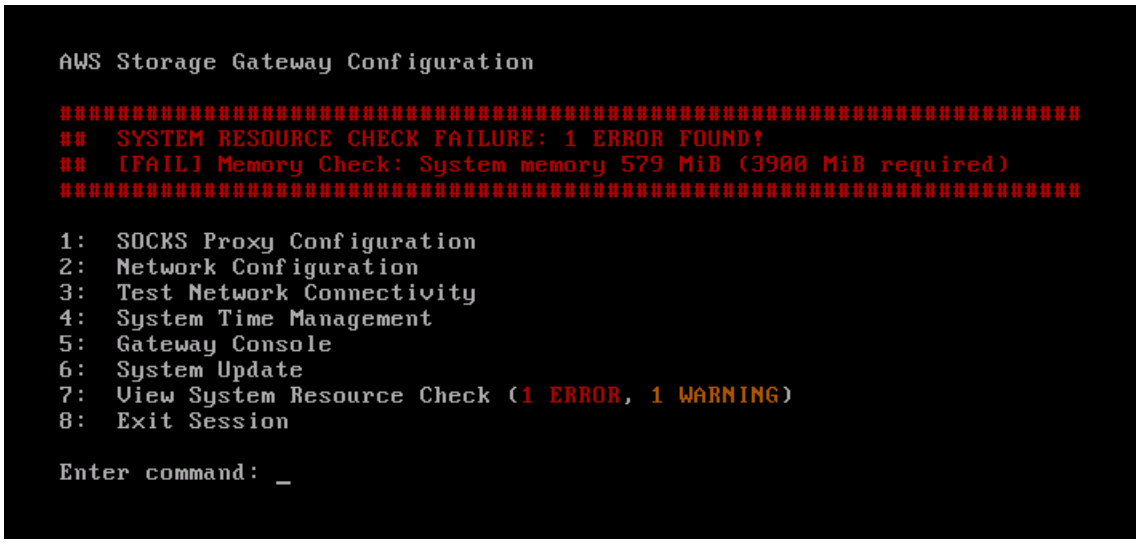
AWS Storage Gateway Configuration
##  SYSTEM RESOURCE CHECK FAILURE: 1 ERROR FOUND!
##  [FAIL] Memory Check: System memory 579 MiB (3900 MiB required)
1: SOCKS Proxy Configuration
2: Network Configuration
3: Test Network Connectivity
4: System Time Management
5: Gateway Console
6: System Update
7: View System Resource Check (1 ERROR, 1 WARNING)
8: Exit Session
Enter command: _

## Configuring Your Gateway for Multiple NICs

If you configure your gateway to use multiple network adapters (NICs), it can be accessed by more than one IP address. You might want to do this in the following situations:

- **Maximizing throughput** – You might want to maximize throughput to a gateway when network adapters are a bottleneck.
- **Application separation** – You might need to separate how your applications write to a gateway's storage volumes. For example, you might choose to have a critical storage application exclusively use one particular adapter defined for your gateway.
- **Network constraints** – Your application environment might require that you keep your iSCSI targets and the initiators that connect to them in an isolated network that is different from the network by which the gateway communicates with AWS.

In a typical multiple-adapter use case, one adapter is configured as the route by which the gateway communicates with AWS (that is, as the default gateway). Except for this one adapter, initiators must be in the same subnet as the adapter that contains the iSCSI targets to which they connect. Otherwise, communication with the intended targets might not be possible. If a target is configured on the same adapter that is used for communication with AWS, then iSCSI traffic for that target and AWS traffic will flow through the same adapter.

The following procedure assumes that your gateway VM already has one network adapter defined and that you are adding a second adapter. The first procedure shows how to add an adapter for VMware ESXi, and the second shows how for Microsoft Hyper-V.

### To configure your gateway to use an additional network adapter (VMware ESXi)

1. Shut down the gateway. For instructions, see To shut down a gateway (p. 127).
2. In the VMware vSphere client, select your gateway VM.

   The VM can remain turned on for this procedure.
3. In the client, right-click the name of your gateway VM, and click **Edit Settings**.

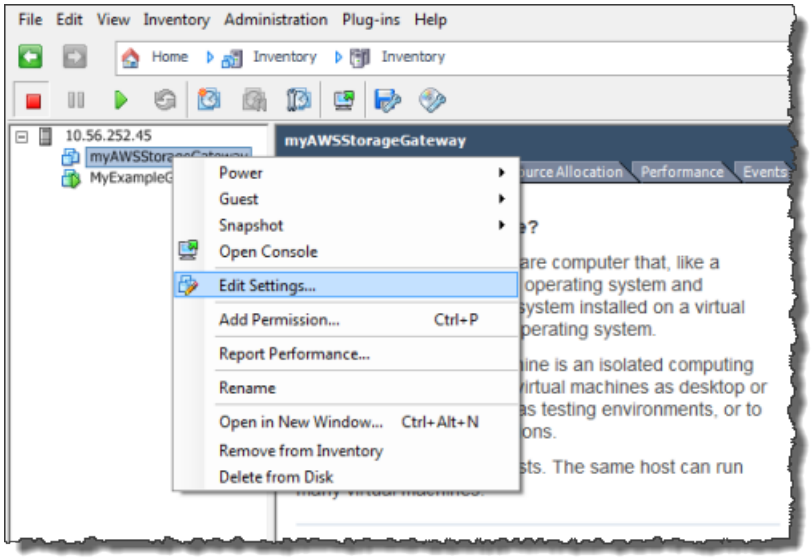


4. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add** to add a device.

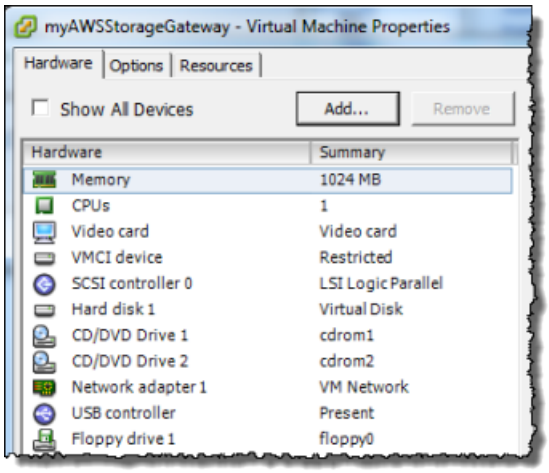


5. Follow the Add Hardware wizard to add a network adapter.

   a. In the **Device Type** pane, click **Ethernet Adapter** to add an adapter, and then click **Next**.

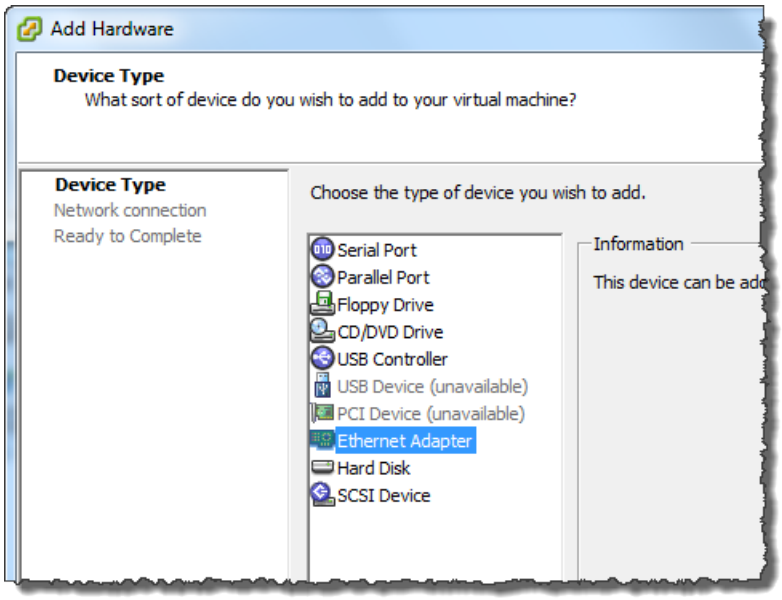


   b. In the **Type** list of the **Network Type** pane, select an adapter type, ensure that **Connect at power on** is selected, and then click **Next**.

      We recommend that you use the E1000 network adapter with AWS Storage Gateway. For more information on the adapter types that might appear in the adapter list, go to Network Adapter Types in the ESXi and vCenter Server Documentation.

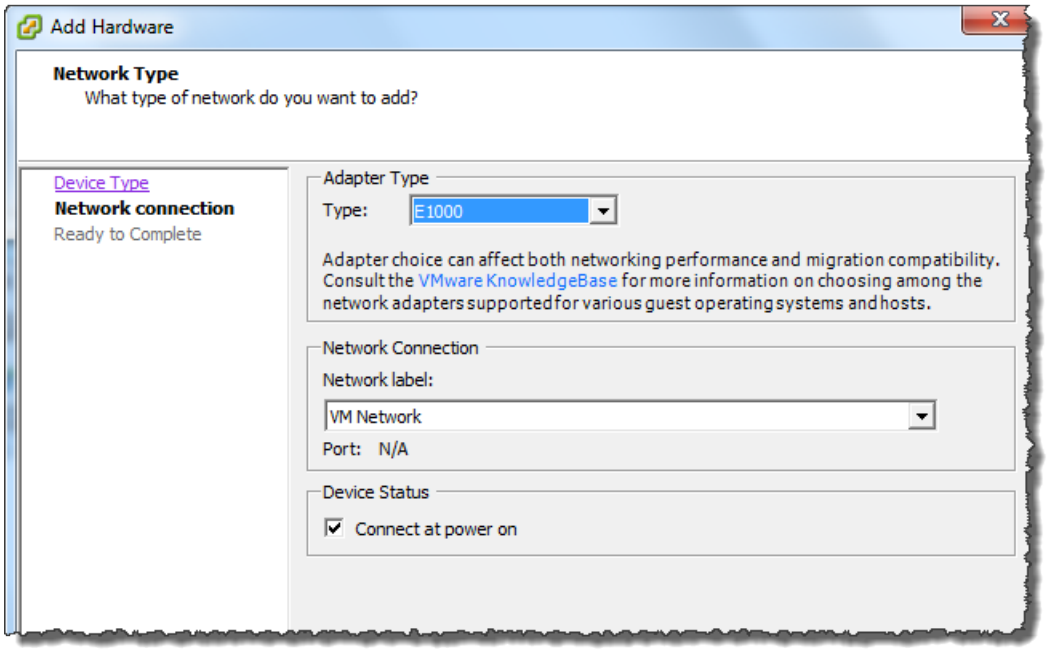
Add Hardware
Network Type
What type of network do you want to add?
Device Type
Network connection
Ready to Complete
Adapter Type
Type:
E1000
Adapter choice can affect both networking performance and migration compatibility. Consult the VMware KnowledgeBase for more information on choosing among the network adapters supported for various guest operating systems and hosts.
Network Connection
Network label:
VM Network
Port: N/A
Device Status
Connect at power on

c. In the **Ready to Complete** pane, review the information, and then click **Finish**.

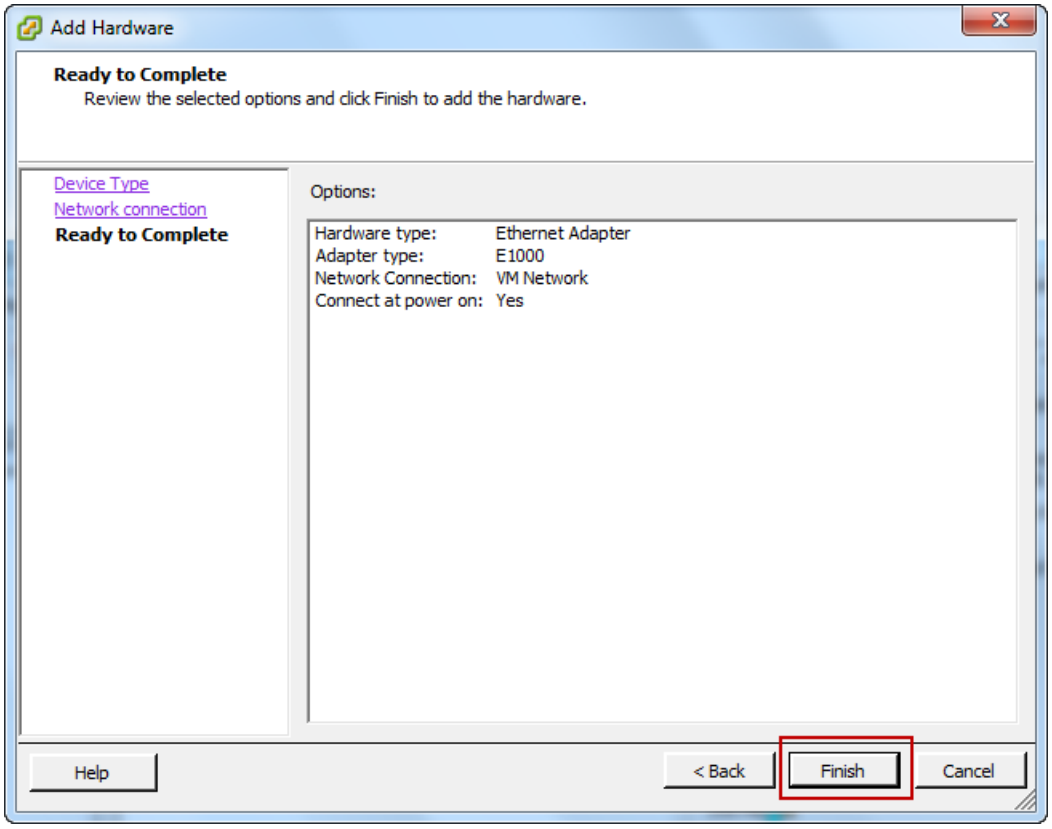


6. Click the **Summary** tab of the VM, and click **View All** next to the **IP Address** box. A **Virtual Machine IP Addresses** window displays all the IP addresses you can use to access the gateway. Confirm that a second IP address is listed for the gateway.

**Note**
It might take several moments for the adapter changes to take effect and the VM summary information to refresh.

The following image is for illustration only. In practice, one of the IP addresses will be the address by which the gateway communicates to AWS and the other will be an address in a different subnet.

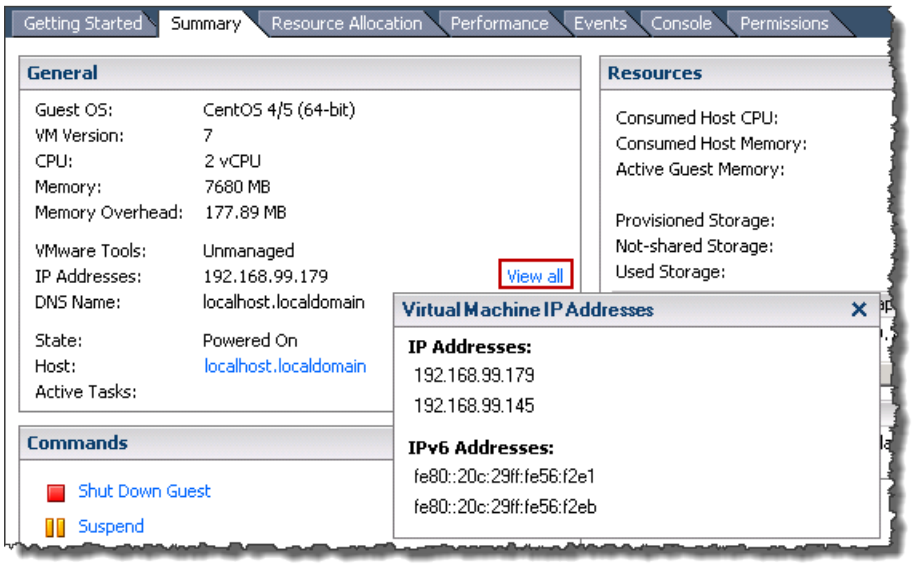


7. On the AWS Storage Gateway console, turn on the gateway. For instructions, see To turn on a gateway (p. 129).
8. In the **Navigation** pane of the AWS Storage Gateway console, click the gateway to which you added the adapter, and then click the **Gateway** tab.

   Confirm that the second IP address is listed in the **IP Addresses** box.

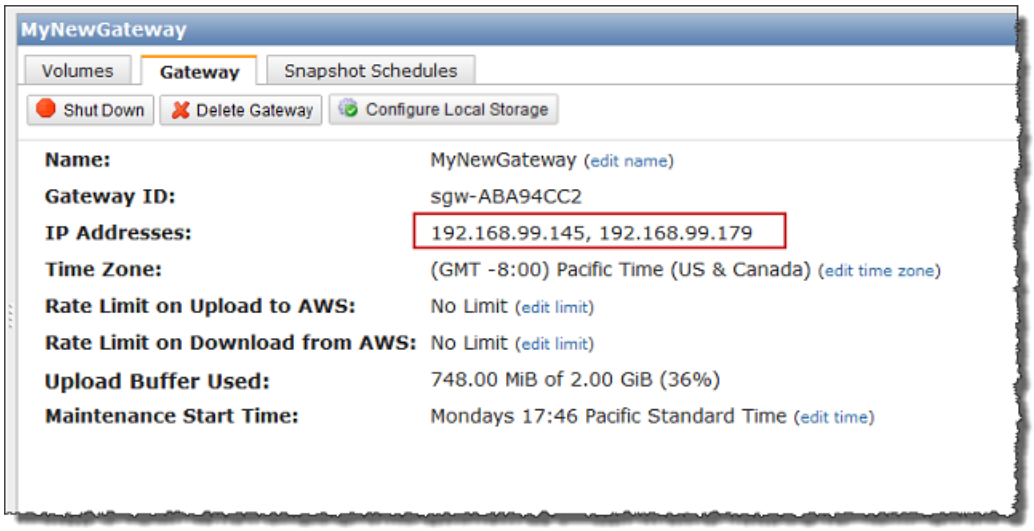
MyNewGateway
Volumes
Gateway
Snapshot Schedules
Shut Down
Delete Gateway
Configure Local Storage
Name: MyNewGateway (edit name)
Gateway ID: sgw-ABA94CC2
IP Addresses: 192.168.99.145, 192.168.99.179
Time Zone: (GMT -8:00) Pacific Time (US & Canada) (edit time zone)
Rate Limit on Upload to AWS: No Limit (edit limit)
Rate Limit on Download from AWS: No Limit (edit limit)
Upload Buffer Used: 748.00 MiB of 2.00 GiB (36%)
Maintenance Start Time: Mondays 17:46 Pacific Standard Time (edit time)

### To configure your gateway to use an additional network adapter (Microsoft Hyper-V)

1. On the AWS Storage Gateway console, turn off the gateway. For instructions, see To shut down a gateway (p. 127).

2. In the Microsoft Hyper-V Manager, select your gateway VM.
3. If the VM it isn't turned off already, right-click the gateway and click **Turn Off**.
4. In the client, right-click the name of your gateway VM and click **Settings**.

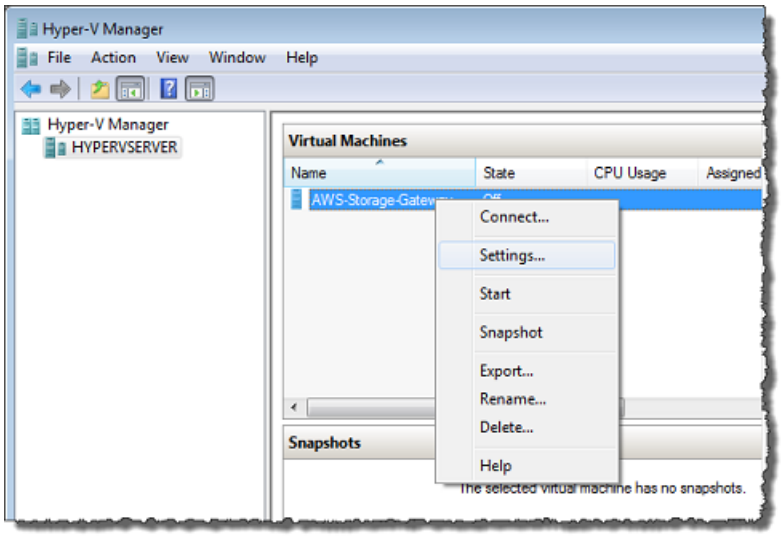


5. In the **Hardware** list in the **Settings** dialog box for the VM, click **Add Hardware**.
6. In the **Add Hardware** pane, click **Network Adapter**, and then click **Add** to add a device.

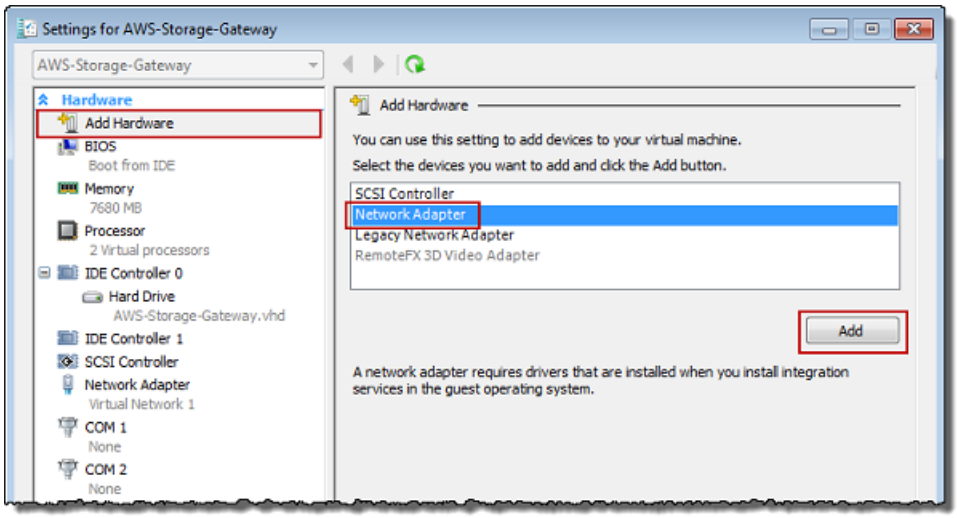


7. Configure the network adapter, and then click **Apply** to apply settings.

   In the following example, **Virtual Network 2** is selected for the new adapter.

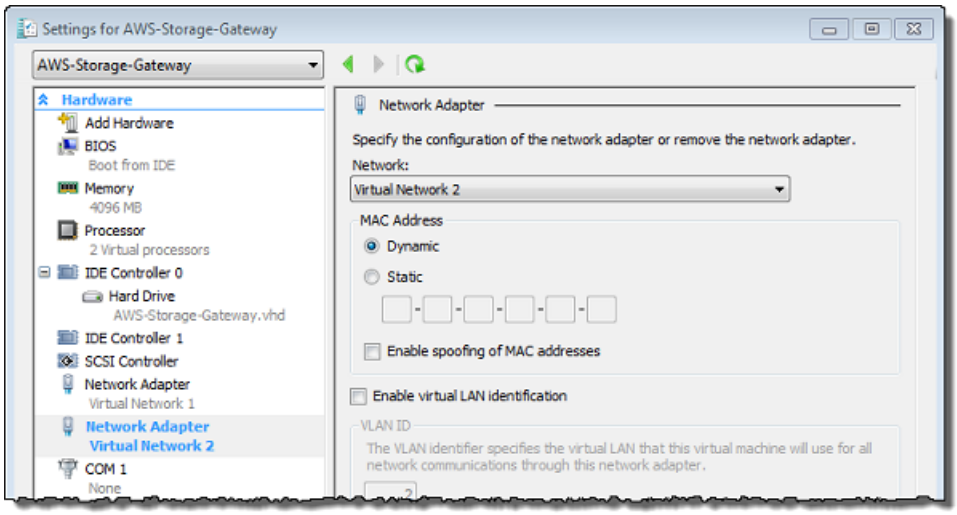


8. In the **Hardware** list of the **Settings** dialog box, confirm that the second adapter was added, and then click **OK**.
9. On the AWS Storage Gateway console, turn on the gateway. For instructions, see To turn on a gateway (p. 129).
10. In the **Navigation** pane of the AWS Storage Gateway console, select the gateway to which you added the adapter, and then click the **Gateway** tab.

    Confirm that the second IP address is listed in the **IP Addresses** field.

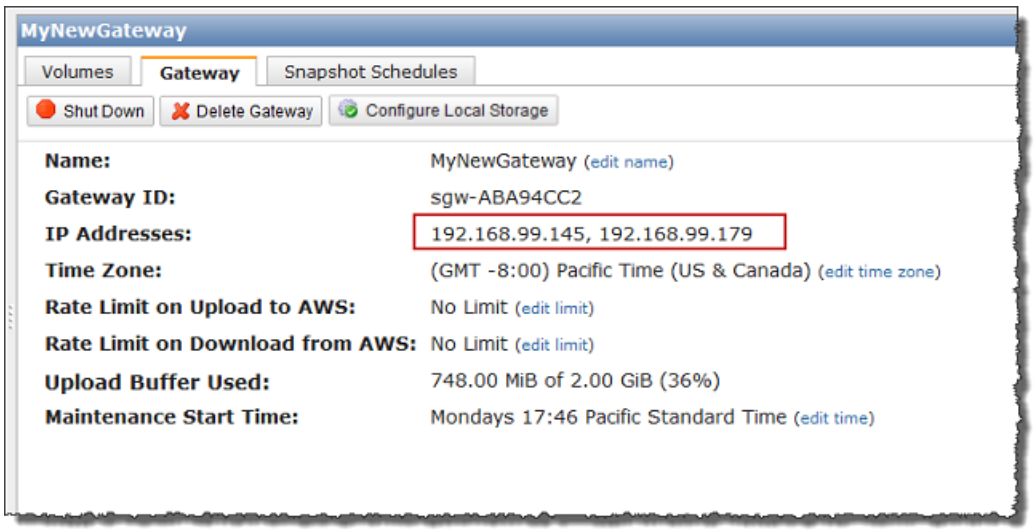
MyNewGateway
Volumes
Gateway
Snapshot Schedules
Shut Down
Delete Gateway
Configure Local Storage
Name: MyNewGateway (edit name)
Gateway ID: sgw-ABA94CC2
IP Addresses: 192.168.99.145, 192.168.99.179
Time Zone: (GMT -8:00) Pacific Time (US & Canada) (edit time zone)
Rate Limit on Upload to AWS: No Limit (edit limit)
Rate Limit on Download from AWS: No Limit (edit limit)
Upload Buffer Used: 748.00 MiB of 2.00 GiB (36%)
Maintenance Start Time: Mondays 17:46 Pacific Standard Time (edit time)

## Creating a Storage Volume on a Gateway with Multiple Network Adapters

If you have defined your gateway to use multiple network adapters, then when you create a storage volume for the gateway you must choose which IP address your storage applications will use to access the volume. Each network adapter defined for a gateway represents one IP address that you can choose. For information about how to add a network adapter to your gateway, see Configuring Your Gateway for Multiple NICs (p. 152).

### To create a storage volume using a specified network adapter

1. In the **Navigation** pane of the AWS Storage Gateway console, select the gateway you want to work with, and click the **Volumes** tab.
2. Click **Create New Volume**.
3. Configure the storage volume. If you have a gateway-cached volume, use the procedure described in Managing Storage Volumes (Gateway-Cached) (p. 97). If you have a gateway-cached volume, use the procedure described in Managing Storage Volumes (Gateway-Stored) (p. 100).
4. Select an IP address to use to access the volume.

   Note that the **Create Storage Volume** dialog box displays a drop-down list for **Host IP**, with one IP address per adapter configured for the gateway VM. If the gateway VM is configured for only one network adapter, the drop-down list is disabled because there is only one IP address.

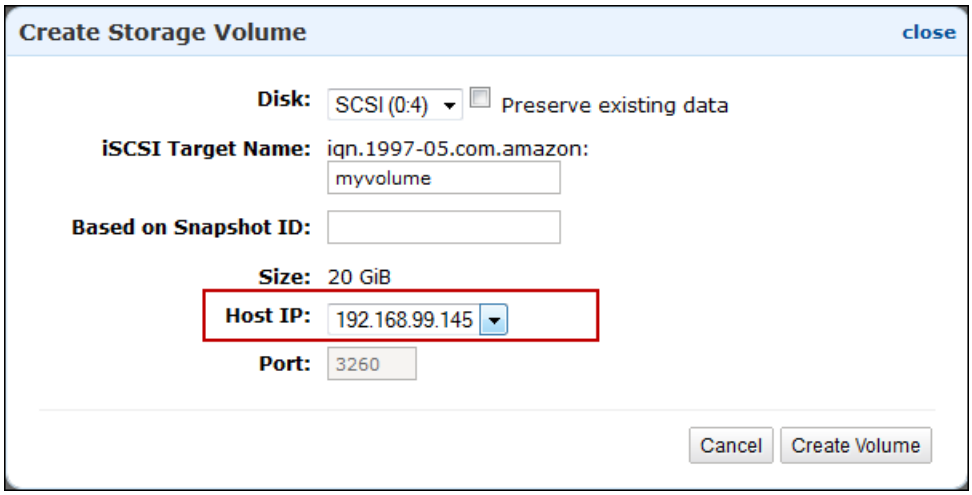


5. Click **Create Volume**.

To create a connection to the storage volume, see Connecting to Storage Volumes (p. 111).

## Optimizing Gateway Performance

This section provides information about how to optimize the performance of your gateway. The guidance is based on adding resources to your gateway and adding resources to your application server.

### Add Resources to Your Gateway

**Use higher-performance disks**

To optimize gateway performance, you can add high performance disks to your gateway. For example, you can add to your gateway serial attached SCSI (SAS) disks and solid-state drives (SSDs), or a redundant array of independent disks, level 5, (RAID 5) backed by four 10,000 RPM 6 Gbps disks.

You can also attach virtual disks to your VM directly from a storage area network (SAN) instead of through the VMware Virtual Machine File System (VMFS) layer or the Microsoft Hyper-V New Technology File System (NTFS). Improved disk performance generally results in better throughput and more input/output operations per second (IOPS). To measure throughput, use the `ReadBytes` and `WriteBytes` metrics with the `Samples` Amazon CloudWatch statistic. For example, the `Samples` statistic of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you the input/output operations per second. As a general rule, when you review these metrics for a gateway, look for low throughput and low IOPS trends to indicate disk-related bottlenecks. For more information about gateway metrics, see Measuring Performance Between Your Gateway and AWS (p. 188).

**Add CPU resources to your gateway host**

The minimum requirement for a gateway host server is four virtual processors. To optimize gateway performance, you should confirm that the four virtual processors that are assigned to the gateway VM are backed by four cores and that you are not oversubscribing the CPUs of the host server. When you add additional CPUs to your gateway host server, you increase the processing capability of the gateway to deal with, in parallel, both storing data from your application to your local storage and uploading this data to Amazon S3. Additional CPUs also help ensure that your gateway gets enough CPU resources when the host is shared with other VMs. Providing enough CPU resources has the general effect of improving throughput.

AWS Storage Gateway supports using eight CPUs in your gateway host server. You can use eight CPUs to significantly improve the performance of your gateway. We recommend the following gateway configuration for your gateway host server:

- Eight CPUs
- 7.5 GiB RAM
- Disk 1 attached to paravirtual controller 1, to be used as the gateway cache as follows:
  - Use RAID 5 backed by four 10,000 RPM 6 Gbps disks with a stripe size of 256K each
  - Use a VMFS5 data store created on the RAID 5
  - Use a virtual disk created on the VMFS5 as the cache for the gateway
- Disk 2 attached to paravirtual controller 2, to be used as the gateway upload buffer as follows:
  - Use RAID 5 backed by four 10,000 RPM 6 Gbps disks with a stripe size of 512K each
  - Use a VMFS5 data store created on the RAID 5
  - Use a virtual disk created on the VMFS5 as the upload buffer for the gateway
- Network adapter 1 configured on VM network 1:
  - Use VM network 1 and add a 1 Gbps network adapter 1 to be used for ingestion
- Network adapter 2 configured on VM network 2:
  - Use VM network 2 and add a 1 Gbps network adapter 2 to be used to connect to AWS

**Change the storage volumes configuration**

If you find that adding more storage volumes to a gateway reduces the throughput to the gateway, consider adding the storage volumes to a separate gateway. In particular, if a storage volume is used for a high-throughput application, consider creating a separate gateway for the high-throughput application. However, as a general rule, you should not use one gateway for all of your high-throughput applications and another gateway for all of your low-throughput applications. To measure your storage volume throughput, use the `ReadBytes` and `WriteBytes` metrics. For more information on these metrics, see Measuring Performance Between Your Application and Gateway (p. 187).

**Back the gateway virtual disks with separate physical disks**

When you provision disks in a gateway-cached volume setup, we strongly recommend that you do not provision local disks for the upload buffer and cache storage that use the same underlying physical storage disk. Similarly, for a gateway-stored volume setup, we strongly recommend that you do not provision local disks for upload buffer and application storage that use the same underlying physical storage disk. For example, for VMware ESXi, the underlying physical storage resources are represented as a data store. When you deploy the gateway VM, you choose a data store on which to store the VM files. When you provision a virtual disk (for example, to use as an upload buffer), you have the option to store the virtual disk in the same data store as the VM or a different data store. If you have more than one data store, then it is strongly recommended that you choose one data store for each type of local storage you are creating. A data store that is backed by only one underlying physical disk, or that is backed by a less-performant RAID configuration such as RAID 1, can lead to poor performance—for example, when used to back both the cache storage and upload buffer in a gateway-cached volume setup.

### Add Resources to Your Application Environment

**Increase the bandwidth between your application server and gateway**
To optimize gateway performance, ensure that the network bandwidth between your application and the gateway can sustain your application needs. You can use the `ReadBytes` and `WriteBytes` metrics of the gateway to measure the total data throughput (for more information on these metrics, see Measuring Performance Between Your Gateway and AWS (p. 188)). For your application, compare the measured throughput with the desired throughput. If the measured throughput is less than the desired throughput, then increasing the bandwidth between your application and gateway can improve performance if the network is the bottleneck. Similarly, you can increase the bandwidth between your VM and your local disks, if they're not direct-attached.

**Add CPU resources to your application environment**
If your application can make use of additional CPU resources, then adding more CPUs can help your application scale its I/O load.

# Maintenance Tasks When Running a Gateway on an Amazon EC2 Instance

In this section, we discuss how you can manage your gateway deployed on Amazon Elastic Compute Cloud (Amazon EC2) as part of a gateway-cached setup. Management tasks you will perform with your gateway include adding and removing Amazon EBS volumes, enabling and disabling AWS Support access to your gateway, configuring security groups, and cleaning up your AWS resources after you are done working with the gateway.

For information about managing a gateway deployed on-premises, see Managing Your Activated Gateway (Gateway-Cached and Gateway-Stored) (p. 96).

**Topics**

- Adding and Removing Amazon EBS Volumes from Your Instance (p. 160)
- Configuring Security Groups for Your Amazon EC2 Gateway Instance (p. 161)
- Cleaning Up Resources After Deleting a Gateway Deployed on Amazon EC2 (p. 162)
- Performing Maintenance Tasks on the Amazon EC2 Gateway Local Console (p. 162)

## Adding and Removing Amazon EBS Volumes from Your Instance

After your gateway is deployed on Amazon Elastic Compute Cloud (Amazon EC2) and activated (see Amazon EC2 Gateway (p. 57)), you might need to configure additional storage to be used as upload buffer and cache storage capacity as your application needs change. Or, you might need to reduce upload buffer or cache storage capacity. In either case, you work with Amazon Elastic Block Store (Amazon EBS) storage, either adding more block storage or reducing it. For more information about Amazon EBS, go to Amazon Elastic Block Store (Amazon EBS) in the *Amazon EC2 User Guide*.

Before you add more storage to the gateway, you should review how to size your upload buffer and cache storage for a gateway-cached setup based on your application needs (see Managing the Upload Buffer and Cache Storage (Gateway-VTL) (p. 304)). Once you configure additional local storage, you work with it just as you would with an on-premises gateway. As an example, see Configuring an Upload Buffer and Cache Storage (Gateway-Cached) (p. 177). For a gateway-cached setup, you can have up to 18 TiB of storage, consisting of up to 2 TiB allocated to an upload buffer and up to 16 TiB allocated to cache storage.

The following tasks demonstrate how to add and remove Amazon EBS volumes from an instance. Although this is a straightforward task for Amazon EC2 instances, you need to take a little extra care when the instance is hosting AWS Storage Gateway. These procedures assume that you already have a deployed and activated gateway.

### To add an Amazon EBS volume to your Amazon EC2–hosted gateway

- To add an Amazon EBS volume to your Amazon EC2–hosted gateway, you create an Amazon EBS volume and then attach it to your Amazon EC2 instance. For information about how to create an Amazon EBS volume, go to Creating a New Amazon EBS Volume in the *Amazon EC2 User Guide.* For information about how to add an Amazon EBS volume to your Amazon EC2–hosted gateway, go to Attaching an Amazon EBS Volume to an Instance in the *Amazon EC2 User Guide.*

### To remove an Amazon EBS volume from your Amazon EC2–hosted gateway

1. Shut down the gateway by following the steps in the Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console (p. 127) procedure.

   **Note**
   Before shutting down the gateway, ensure that no application is writing data to it and that no snapshots are in progress. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the AWS Storage Gateway console. For more information, see Editing a Snapshot Schedule (p. 210).
2. In the **Navigation** pane of the Amazon EC2 console, click **Instances**.
3. Find the instance running the gateway.
4. Among the block devices attached to the instance, find the device you want to remove.

   Note that the root device and the swap device are listed as **X** and **Y** and should not be removed.
5. In the **Navigation** pane of the Amazon EC2 console, click **Volumes**.
6. Find the volume you want to remove from the gateway, right-click on it, and select **Force Detach**.

## Configuring Security Groups for Your Amazon EC2 Gateway Instance

A security group controls traffic to your Amazon EC2 gateway instance. When you create an instance from the Amazon Machine Image (AMI) for AWS Storage Gateway from AWS Marketplace, you have two choices for launching the instance. In Amazon EC2 Gateway (p. 57), we show how to launch the instance by using the **Launch with an EC2 Console** feature of AWS Marketplace. Using this feature is the recommended approach. You can also launch an instance by using the **1-Click Launch** feature in AWS Marketplace. In this case, the auto-generated security group `AWS Storage Gateway-1-0-AutogenByAWSMP-` is created. This security group has the correct port 80 rule to allow you to activate your gateway. For more information about security groups, go to Security Group Concepts in the *Amazon EC2 User Guide.*

Regardless of the security group that you use, we recommend that:

- The security group should not allow incoming connections from the outside Internet. It should allow only instances within the appliance's security group to communicate with the appliance. If you need to allow instances to connect to the appliance from outside the appliance's security group, we recommend that you allow connections only on ports 3260 (for iSCSI) and 80 (for activation).
- You allow port 22 access only if you are using AWS Support for troubleshooting purposes. For more information, see Enabling AWS Support to Access a Gateway Hosted on an Amazon EC2 Instance (p. 232).
- If you want to activate your appliance from a host outside the appliance's security group, you will need to allow incoming connections on port 80 from the IP address of that host. If you cannot determine the activating host's IP address, you can open port 80, activate your gateway, and then close access on port 80 after completing activation.

If you are using an Amazon EC2 instance as an initiator, that is to connect to the iSCSI targets on the gateway you deployed on Amazon EC2, then you have two options. You can put the initiator instance in

the same security group as the gateway, or you can configure access so the initiator can communicate with the gateway.

## Cleaning Up Resources After Deleting a Gateway Deployed on Amazon EC2

If you are done using a gateway you deployed on Amazon EC2, we recommend that you clean up the AWS resources that were used for the gateway, specifically the Amazon EC2 instance and any Amazon EBS volumes. Doing so helps avoid unintended usage charges. We suggest that you take the following actions for deleting your gateway and cleaning up its resources:

- In the AWS Storage Gateway console, delete the gateway as shown in Deleting a Gateway Using the AWS Storage Gateway Console (p. 136).
- In the Amazon EC2 console, stop the instance if you plan on using the instance again. Terminate the instance if you do not plan on using the instance again. Before terminating the instance, note the block devices that are attached to the instance and these devices' identifiers if you plan on deleting volumes.
- In the Amazon EC2 console, remove any Amazon EBS volumes that are attached to the instance if you do not plan on using them again.

## Performing Maintenance Tasks on the Amazon EC2 Gateway Local Console

Some maintenance tasks require that you log in to the local console when running a gateway deployed on an Amazon EC2 instance. This section describes how to log in to the local console and perform maintenance tasks.

**Topics**


- Logging In to Your Amazon EC2 Gateway Local Console (p. 162)
- Routing Your Gateway Deployed on EC2 Through a Proxy (p. 163)
- Testing Your Gateway Connectivity to the Internet (p. 164)
- Running Storage Gateway Commands on the Local Console (p. 166)
- Updating Your System (p. 167)
- Viewing Your Gateway System Resource Errors (p. 168)


### Logging In to Your Amazon EC2 Gateway Local Console

You can connect to your Amazon EC2 instance by using a Secure Shell (SSH) client. For detailed information, go to Connect to Your Instance in the *Amazon EC2 User Guide*. To connect this way, you will need the SSH key pair you specified when you launched the instance. For information about Amazon EC2 key pairs, go to Amazon EC2 Key Pairs in the *Amazon EC2 User Guide.*

#### To log in to the gateway local console

1. Log in to your local console. If you are connecting from a Windows computer, log in as *sguser*.
2. After you log in, you will see the **AWS Storage Gateway Configuration** main menu, as shown in the following screenshot.

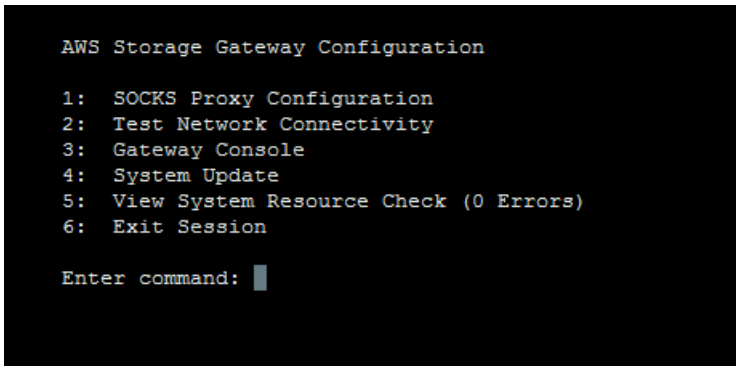


| To... | See... |
|---|---|
| Configure a SOCKS proxy for your gateway | Routing Your On-Premises Gateway Through a Proxy (p. 141) |
| Test network connectivity | Testing Your Gateway Connection to the Internet (p. 145) |
| Run Storage Gateway console commands | Running Storage Gateway Commands on the Local Console (p. 148) |
| Update your system | Updating Your System (p. 149) |

To exit the configuration session, type **5** to exit the menu.

## Routing Your Gateway Deployed on EC2 Through a Proxy

AWS Storage Gateway supports the configuration of a Socket Secure version 5 (SOCKS5) proxy between your gateway deployed on Amazon EC2 and AWS.

> **Note**
> The only proxy configuration AWS Storage Gateway supports is SOCKS5.

If your gateway must use a proxy server to communicate to the Internet, then you need to configure the SOCKS proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy. After you do so, AWS Storage Gateway will route all HyperText Transfer Protocol Secure (HTTPS) traffic through your proxy server.

### To route your gateway Internet traffic through a local proxy server

1. Log in to your gateway's local console. For instructions, see Logging In to Your Amazon EC2 Gateway Local Console (p. 162).
2. On the **AWS Storage Gateway Configuration** main menu, type **1** to begin configuring the SOCKS proxy.

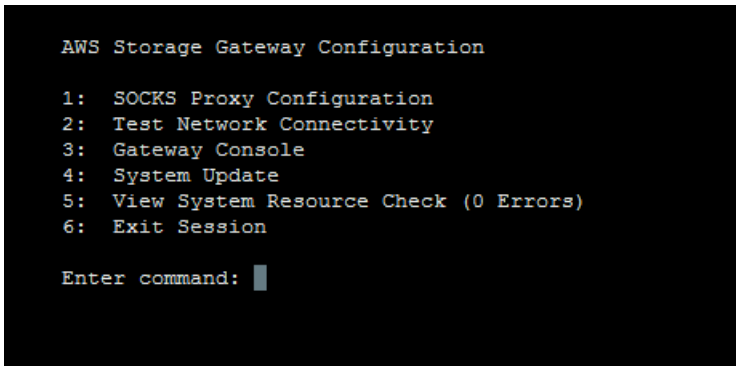


3. Choose one of the following options in the **AWS Storage Gateway SOCKS Proxy Configuration** menu:

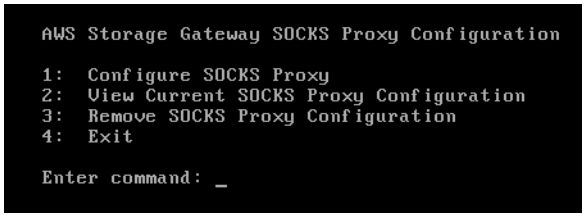
AWS Storage Gateway SOCKS Proxy Configuration
1: Configure SOCKS Proxy
2: View Current SOCKS Proxy Configuration
3: Remove SOCKS Proxy Configuration
4: Exit
Enter command: _

| To... | Do this... |
|---|---|
| Configure a SOCKS proxy | Type **1**.<br><br>You will need to supply a host name and port to complete configuration. |
| View the current SOCKS proxy configuration | Type **2**.<br><br>If a SOCKS proxy is not configured, the message `SOCKS Proxy not configured` is displayed. If a SOCKS proxy is configured, the host name and port of the proxy are displayed. |
| Remove a SOCKS proxy configuration | Type **3**.<br><br>The message `SOCKS Proxy Configuration Removed` is displayed. |
| Exit this menu and return to the previous menu | Type **4**. |

## Testing Your Gateway Connectivity to the Internet

You can use your gateway's local console to test your Internet connection. This test can be useful when you are troubleshooting network issues with your gateway. The network connectivity test does not take into account any SOCKS proxy you might have configured.

### To test your gateway's connection to the Internet

1. Log in to your gateway's local console. For instructions, see Logging In to Your Amazon EC2 Gateway Local Console (p. 162).

2. In the **AWS Storage Gateway Configuration** main menu, type **2** to begin testing network connectivity.

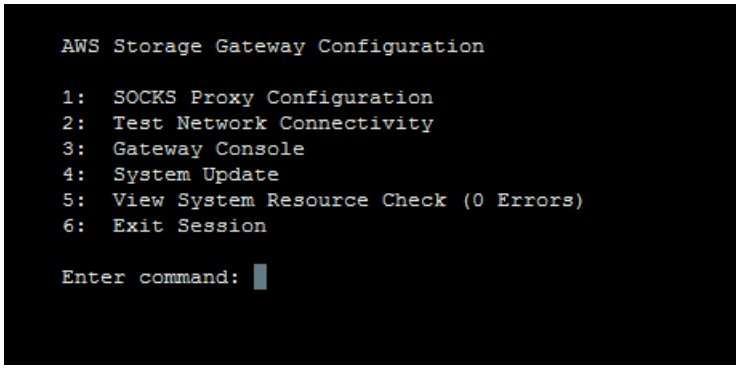


The console displays the available regions.

3. Select the region you want to test, as shown in the following screenshot.

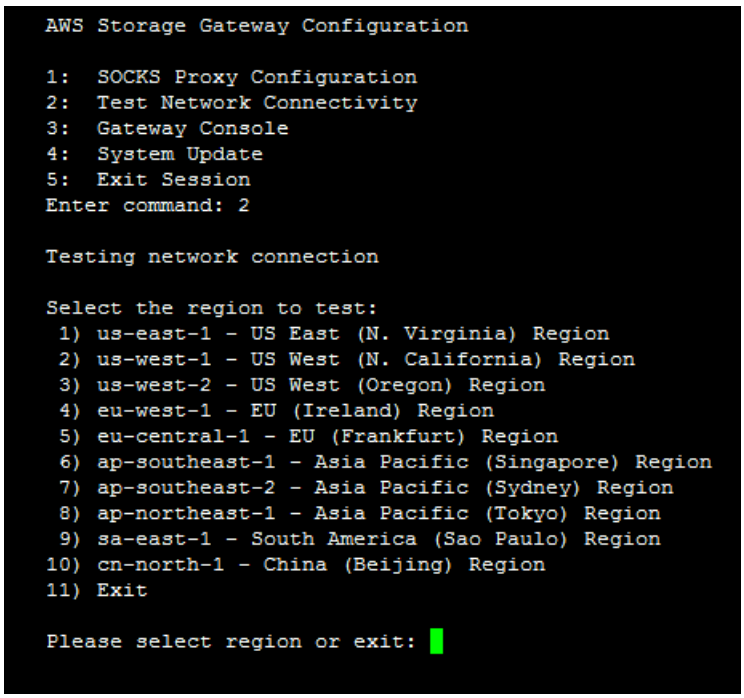
AWS Storage Gateway Configuration
1: SOCKS Proxy Configuration
2: Test Network Connectivity
3: Gateway Console
4: System Update
5: Exit Session
Enter command: 2
Testing network connection
Select the region to test:
1) us-east-1 - US East (N. Virginia) Region
2) us-west-1 - US West (N. California) Region
3) us-west-2 - US West (Oregon) Region
4) eu-west-1 - EU (Ireland) Region
5) eu-central-1 - EU (Frankfurt) Region
6) ap-southeast-1 - Asia Pacific (Singapore) Region
7) ap-southeast-2 - Asia Pacific (Sydney) Region
8) ap-northeast-1 - Asia Pacific (Tokyo) Region
9) sa-east-1 - South America (Sao Paulo) Region
10) cn-north-1 - China (Beijing) Region
11) Exit
Please select region or exit:

Each endpoint in the region you select displays either a PASSED or FAILED message, as shown following.

| Message | Description |
|---|---|
| [ PASSED ] | AWS Storage Gateway has Internet connectivity. |
| [ FAILED ] | AWS Storage Gateway does not have Internet connectivity. |

## Running Storage Gateway Commands on the Local Console

The AWS Storage Gateway console helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the console commands, you can perform maintenance tasks such as saving routing tables or connecting to AWS Support.

### To run a configuration or diagnostic command

1. Log in to your gateway's local console. For instructions, see Logging In to Your Amazon EC2 Gateway Local Console (p. 162).
2. In the **AWS Storage Gateway Configuration** main menu, type **3** for **Gateway Console**.

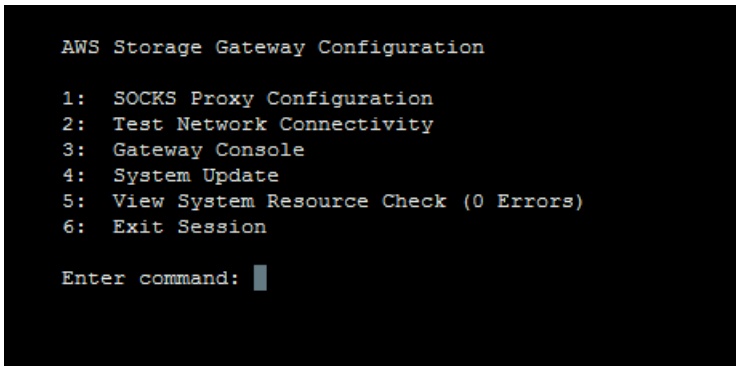


3. In the Storage Gateway console, type **h**, and then press the **Return** key.

The console displays the **Available Commands** menu with the available commands and after the menu a **Gateway Console** prompt, as shown in the following screenshot.

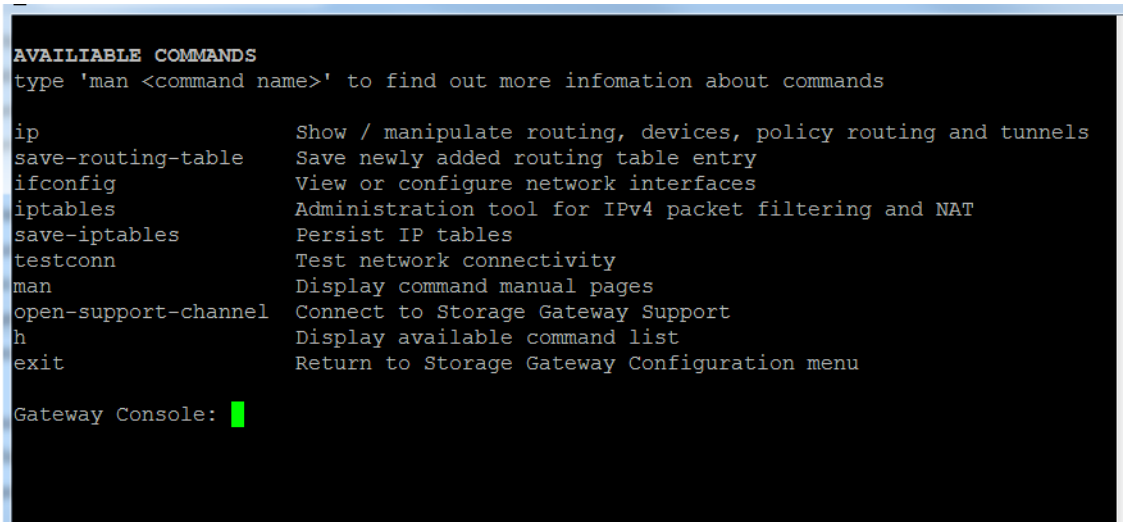
AVAILIABLE COMMANDS
type 'man <command name>' to find out more infomation about commands
ip Show / manipulate routing, devices, policy routing and tunnels
save-routing-table Save newly added routing table entry
ifconfig View or configure network interfaces
iptables Administration tool for IPv4 packet filtering and NAT
save-iptables Persist IP tables
testconn Test network connectivity
man Display command manual pages
open-support-channel Connect to Storage Gateway Support
h Display available command list
exit Return to Storage Gateway Configuration menu
Gateway Console:

4. To learn about a command, type **man** + ***command name*** at the **Gateway Console** prompt.

## Updating Your System

You can use the local console to update the software on your system.

**Warning**
Applying system updates when the gateway is turned on can cause volumes to move to PASS THROUGH state. Shut down the gateway before applying system updates.

### To update your systems

1. Log in to your gateway's local console. For instructions, see Logging In to Your Amazon EC2 Gateway Local Console (p. 162).
2. On the **AWS Storage Gateway Configuration** main menu, type **4** to begin updating.

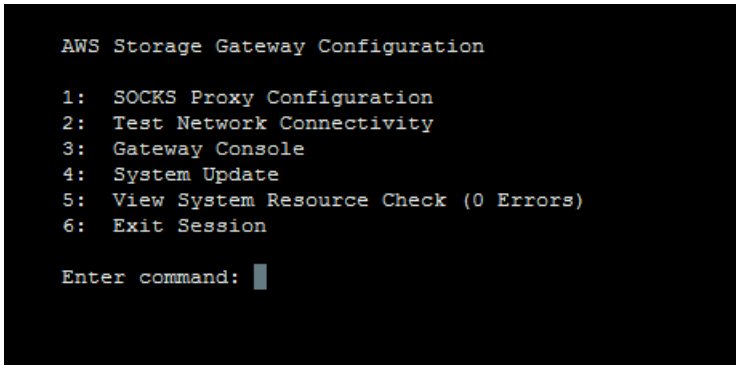


The console displays the **Apply System Updates** menu, as shown in the following screenshot.

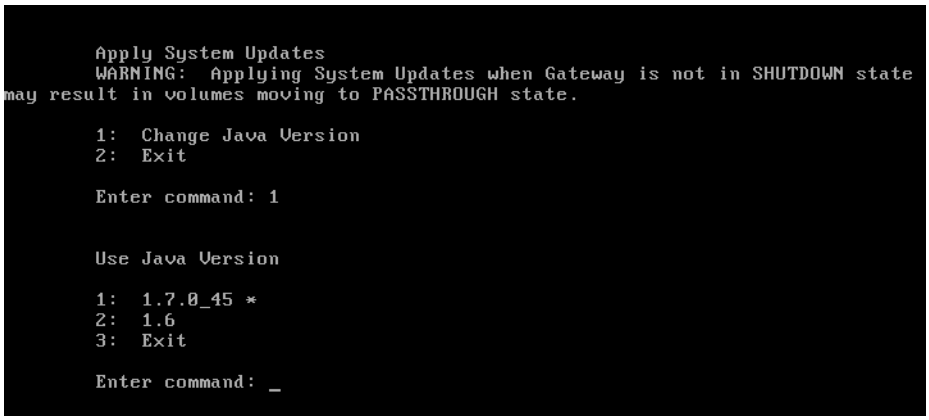
Apply System Updates
WARNING: Applying System Updates when Gateway is not in SHUTDOWN state may result in volumes moving to PASSTHROUGH state.
1: Change Java Version
2: Exit
Enter command: 1
Use Java Version
1: 1.7.0_45 *
2: 1.6
3: Exit
Enter command: _

3. Type one of the following options to use a new version of Java or to exit.

| To... | Do this... |
|---|---|
| To use Java Version 1.7.0_45 * | Type **1**. |
| To use Java Version 1.6 | Type **2**. |
| Exit to the previous menu | Type **3**. |

## Viewing Your Gateway System Resource Errors

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM and determines whether these system resources are sufficient for your gateway to function properly. You can view the results of this check on the gateway's local console.

### To view the status of a system resource check

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 137).
2. In the **AWS Storage Gateway Configuration** main menu, type **5** to view the results of a system resource check.

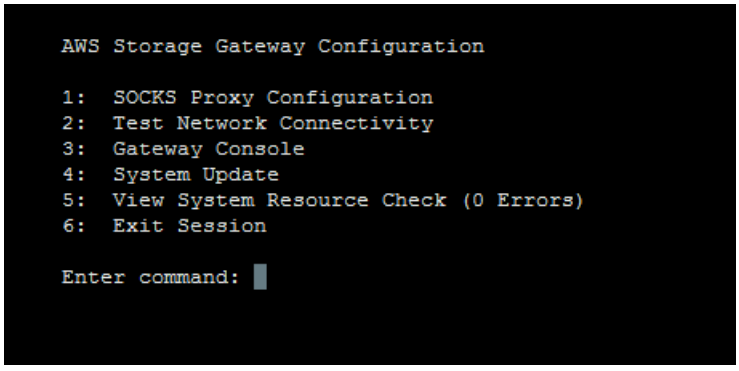


The console displays an **[OK]**, **[WARNING]**, or **[FAIL]** message for each resource as described in the table following.

| Message | Description |
|---|---|
| **[OK]** | The resource has passed the system resource check. |
| **[WARNING]** | The resource does not meet the recommended requirements, but your gateway will continue to function. AWS Storage Gateway displays a message that describes the results of the resource check. |
| **[FAIL]** | The resource does not meet the minimum requirements. Your gateway might not function properly. AWS Storage Gateway displays a message that describes the results of the resource check. |

The console also displays the number of errors and warnings next to the resource check menu option.

The following screenshot shows a **[FAIL]** message and the accompanying error message.

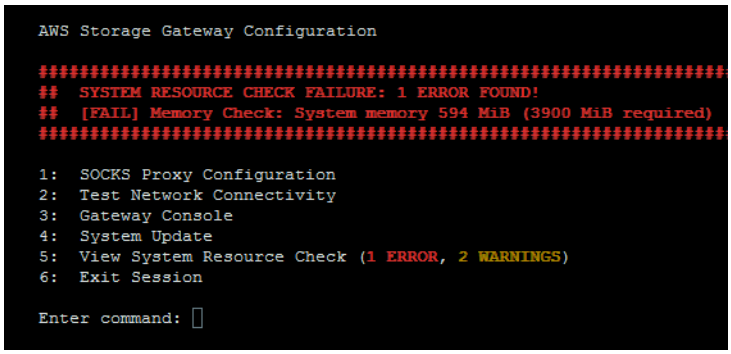
AWS Storage Gateway Configuration
##  SYSTEM RESOURCE CHECK FAILURE: 1 ERROR FOUND!
##  [FAIL] Memory Check: System memory 594 MiB (3900 MiB required)
1:  SOCKS Proxy Configuration
2:  Test Network Connectivity
3:  Gateway Console
4:  System Update
5:  View System Resource Check (1 ERROR, 2 WARNINGS)
6:  Exit Session
Enter command:

# Managing the Upload Buffer and Cache Storage (Gateway-Cached)

The gateway VM can use the local disks that you allocate on-premises for the following purposes.

**Upload buffer**
The upload buffer provides a staging area for the data before the gateway uploads the data to Amazon S3. Your gateway uploads this buffer data over an encrypted Secure Sockets Layer (SSL) connection to AWS.

**Cache storage**
The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer. When your application performs I/O on a volume, the gateway saves the data to the cache storage for low-latency access. When your application requests data from a volume, the gateway first checks the cache storage for the data before downloading the data from AWS.

**Topics**

- Managing the Upload Buffer (Gateway-Cached) (p. 170)
- Managing Cache Storage (Gateway-Cached) (p. 172)
- Sizing Cache Storage (p. 173)
- Adding and Removing Upload Buffer Capacity (Gateway-Cached) (p. 173)

- Adding Cache Storage (Gateway-Cached) (p. 177)
- Configuring an Upload Buffer and Cache Storage (Gateway-Cached) (p. 177)

**Note**
When you provision disks, we strongly recommend that you do not provision local disks for the upload buffer and cache storage that use the same underlying physical storage resource (that is, the same disk). Underlying physical storage resources are represented as a data store in VMware. When you deploy the gateway VM, you choose a data store on which to store the VM files. When you provision a local disk (for example, to use as cache storage or upload buffer), you have the option to store the virtual disk in the same data store as the VM or a different data store. If you have more than one data store, we strongly recommend that you choose one data store for the cache storage and another for the upload buffer. A data store that is backed by only one underlying physical disk, or that is backed by a less-performant RAID configuration such as RAID 1, can lead to poor performance in some situations when used to back both the cache storage and upload buffer.

## Managing the Upload Buffer (Gateway-Cached)

The following diagram highlights the upload buffer in the larger picture of the gateway-cached architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 2).

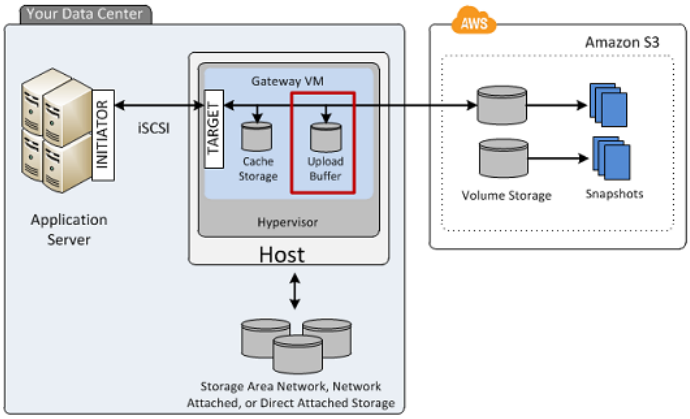


The amount of upload buffer space your gateway requires depends on several factors, such as the rate of incoming data to the storage volumes, the rate of outgoing data to AWS, and your network bandwidth. If your applications continue to write data at a fast rate to your storage volumes, and network throughput is not sufficient for the gateway to upload data to AWS, then eventually your upload buffer will be filled with data waiting to be uploaded to AWS. Here is some guidance you can follow to avoid this situation:

- **Use the sizing formula** – As your application needs change, you should periodically review the recommended formula for sizing the upload buffer. For more information, see Sizing the Upload Buffer (p. 171).
- **Use Amazon CloudWatch metrics** – You can proactively avoid the upload buffer from filling up by monitoring the percentage of upload buffer space your gateway is using over time. Amazon CloudWatch provides metrics such as the `UploadBufferPercentUsed` metric for monitoring your gateway's upload buffer usage (for more information, see Monitoring the Upload Buffer (p. 191)). You can set an alarm to trigger a notification to you when upload buffer usage exceeds a threshold. If your upload buffer is getting filled close to capacity, consider adding more buffer capacity to the gateway. For a full list of AWS Storage Gateway metrics, see Understanding AWS Storage Gateway Metrics (p. 195).

  You can see the current percentage of upload buffer usage on the **Gateway** tab of the AWS Storage Gateway console.

- **Monitor volume status** – A volume status can indicate an issue with the upload buffer. For example, if the upload buffer capacity is reached, the storage volumes affected go into a PASS THROUGH (p. 108) status. Although your applications can continue to operate, writing and reading data to and from your storage volume, snapshots are not taken during this time.
- **Optimize your environment** – If the speed of incoming write operations is too high compared to the outgoing network bandwidth, then the gateway might never catch up, no matter how much upload buffer capacity you provision. In this case, consider optimizing your gateway for better performance. For more information, see Optimizing Gateway Performance (p. 158).

## Sizing the Upload Buffer

You can determine the size of your upload buffer by using an upload buffer formula. We strongly recommend that you allocate at least 150 GiB of upload buffer. If the formula returns a value less than 150 GiB, use 150 GiB as the amount you allocate to the upload buffer. You can configure up to 2 TiB of upload buffer capacity for each gateway.

> **Note**
> When the upload buffer reaches its capacity, your volume goes to PASS THROUGH status. In this status, new data that your application writes is persisted locally but not uploaded to AWS immediately. Thus, you cannot take new snapshots. When the upload buffer capacity frees up, the volume goes through BOOTSTRAPPING status. In this status, any new data that was persisted locally is uploaded to AWS. Finally, the volume returns to ACTIVE status. AWS Storage Gateway then resumes normal synchronization of the data stored locally with the copy stored in AWS, and you can start taking new snapshots. For more information about volume status, see Understanding Storage Volume Status (p. 105).

To estimate the amount of upload buffer, you can determine the expected incoming and outgoing data rates and plug them into the following formula.

**Rate of incoming data**
This rate refers to the application throughput, the rate at which your on-premises applications write data to your gateway over some period of time.

**Rate of outgoing data**
This rate refers to the network throughput, the rate at which your gateway is able to upload data to AWS. This rate depends on your network speed, utilization, and whether you've enabled bandwidth throttling. This rate should be adjusted for compression. When uploading data to AWS, the gateway applies data compression where possible. For example, if your application data is text-only, you might get an effective compression ratio of about 2:1. However, if you are writing videos, the gateway might not be able to achieve any data compression and might require more upload buffer for the gateway.

If your incoming rate is higher than the outgoing rate, or if the formula returns a value less than 150 GiB, we strongly recommend that you allocate at least 150 GiB of upload buffer space.

For example, assume that your business applications will write text data to your gateway at a rate of 40 MB per second for 12 hours a day and your network throughput is 12 MB per second. Assuming a compression factor of 2:1 for the text data, you need to allocate approximately 690 GiB of space for the upload buffer.

```
((40 MB/sec) - (12 MB/sec * 2)) * (12 hours * 3600 seconds/hour) = 691200
megabytes
```

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway as upload buffer space. Add more upload buffer space as needed using the AWS Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see Monitoring the Upload Buffer (p. 355).

If you decide that you need to change your upload buffer capacity, take one of the following actions.

| To... | Do This... |
|---|---|
| Add more upload buffer capacity to your gateway. | Follow the steps in Adding Upload Buffer Capacity (p. 173). |
| Remove a disk allocated as upload buffer space. | Follow the steps in Removing Upload Buffer Capacity (p. 173). |

## Managing Cache Storage (Gateway-Cached)

The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer.

The following diagram highlights the cache storage in the larger picture of the AWS Storage Gateway gateway-cached architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 2).

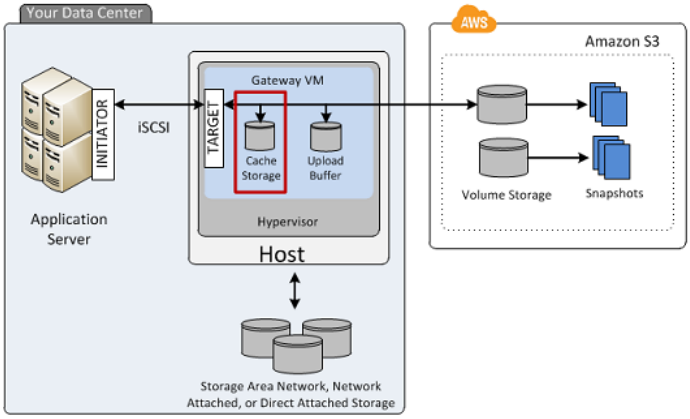


The amount of cache storage your gateway requires depends on how much of your application data you want to provide low-latency access to. The cache storage must be at least the size of the upload buffer. This approach helps ensure that the cache storage is large enough to persistently hold all data that has not yet been uploaded to Amazon S3. When your cache storage has filled up with dirty data (that is, data that has not been uploaded to AWS), application write operations to your storage volume are blocked until more cache storage becomes available. However, application read operations from the storage volume are still allowed. Here is some guidance you can follow to avoid this situation:

- **Use the sizing formula.** – As your application needs change, you should periodically review the recommended formula for sizing cache storage. For more information, see Sizing Cache Storage (p. 173).
- **Use Amazon CloudWatch metrics.** – You can proactively avoid filling up cache storage with dirty data by monitoring how cache storage is being used—particularly, by reviewing cache misses. Amazon CloudWatch provides usage metrics such as the `CachePercentDirty` and `CacheHitPercent` metrics for monitoring how much of the gateway's cache storage has not been uploaded to Amazon S3. You can set an alarm to trigger a notification to you when the percentage of the cache that is dirty exceeds

a threshold or the cache hit percentage falls below a threshold. Both of these can indicate that the cache storage size is not adequate for the gateway. For a full list of AWS Storage Gateway metrics, see Understanding AWS Storage Gateway Metrics (p. 195).

## Sizing Cache Storage

The gateway uses the cache storage to provide low-latency access to your recently accessed data. The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer. Generally speaking, you size the cache storage at 1.1 times the upload buffer size. For more information about how to estimate your cache storage size, see Sizing the Upload Buffer (p. 171).

You can initially use this approximation to provision disks for the cache storage. You can then use Amazon CloudWatch operational metrics to monitor the cache storage usage and provision more storage as needed using the console. For information on using the metrics and setting up alarms, see Monitoring Cache Storage (p. 194).

If you decide that you need to increase your gateway's cache storage capacity, follow the steps in Adding Cache Storage (Gateway-Cached) (p. 177).

## Adding and Removing Upload Buffer Capacity (Gateway-Cached)

After you configure your initial gateway cache storage (see Configuring an Upload Buffer and Cache Storage (Gateway-Cached) (p. 177)), you can configure additional upload buffer capacity as your application needs change. To learn more about how to size your upload buffer based on your application needs, see Sizing the Upload Buffer (p. 171).

You can add more buffer capacity to your gateway without interrupting existing gateway functions. Note that when you add more upload buffer capacity, you do so with the gateway VM turned on. However, when you reduce the amount of upload buffer capacity, you must first turn off the VM.

### Adding Upload Buffer Capacity

As your application needs change and you add more storage volume capacity, you might need to increase the gateway's upload buffer capacity. You can add more buffer capacity by using the AWS Storage Gateway console or the AWS Storage Gateway API:

- For information on adding buffer capacity with the console, see To configure local disks as upload buffer and cache storage (p. 177). This procedure assumes that your activated gateway has at least one local disk available on its VM that you can allocate as an upload buffer to the gateway.
- For information on adding buffer capacity with the API, see AddUploadBuffer.

### Removing Upload Buffer Capacity

As your application needs change and you change the storage volume configuration for a gateway, you might need to decrease the gateway's upload buffer capacity. Or, a local disk allocated as upload buffer space might fail and you might need to remove that disk from your upload buffer and assign a new local disk. In both cases, you can remove buffer capacity using the AWS Storage Gateway console. The following procedure assumes that your activated gateway has at least one local disk allocated as an upload buffer for the gateway. In the procedure, you start on the AWS Storage Gateway console, leave the console and use the VMware vSphere client or the Microsoft Hyper-V Manager to remove the disk, and then return to the console.

#### To find the ID of a disk allocated as an upload buffer

1. On the **Gateway** tab of the AWS Storage Gateway console, click **Configure Local Storage**.

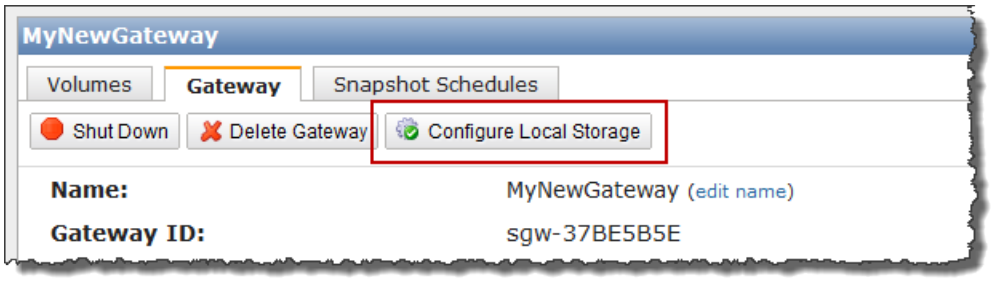


2. In the **Configure Your Activated Gateway** dialog box, note the value of the virtual device node for the local disk to be removed. You can find the node value in the **Local Disks** column. For example, in the following dialog box, the device node value **SCSI (0:2)** is highlighted.

   You use the disk's virtual device node in the vSphere client to help ensure that you remove the correct disk.

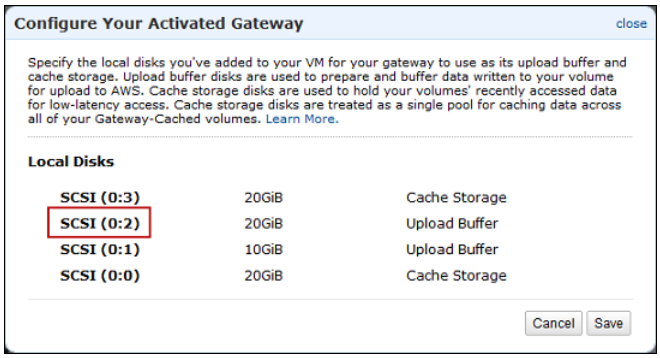
Configure Your Activated Gateway
close
Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage. Upload buffer disks are used to prepare and buffer data written to your volume for upload to AWS. Cache storage disks are used to hold your volumes' recently accessed data for low-latency access. Cache storage disks are treated as a single pool for caching data across all of your Gateway-Cached volumes. Learn More.
Local Disks
SCSI (0:3) 20GiB Cache Storage
SCSI (0:2) 20GiB Upload Buffer
SCSI (0:1) 10GiB Upload Buffer
SCSI (0:0) 20GiB Cache Storage
Cancel
Save

3. Shut down the gateway by following the steps in the Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console (p. 127) procedure.

   **Note**
   Before shutting down the gateway, ensure that no application is writing data to it and that no snapshots are in progress. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the AWS Storage Gateway console. For more information, see Editing a Snapshot Schedule (p. 210).

4. To remove the underlying local disk, do one of the following procedures.

| For a Gateway Hosted In... | Do This... |
|---|---|
| VMware ESXi | Follow the steps in To remove an underlying disk allocated for the upload buffer (VMware ESXi) (p. 175). |
| Microsoft Hyper-V | Follow the steps in To remove an underlying disk allocated for the upload buffer (Microsoft Hyper-V) (p. 176). |

5. On the AWS Storage Gateway console, turn on the gateway.

   **Important**
   After removing a disk used as an upload buffer, you must turn the gateway back on before adding new disks to the VM.

6. On the **Volumes** tab of the AWS Storage Gateway console, check that all storage volumes have a status of AVAILABLE (p. 108).

After a gateway restart, a storage volume might go through the PASS THROUGH (p. 108) and BOOTSTRAPPING (p. 108) states as the gateway adjusts to the upload buffer disk that you removed. A storage volume that passes through these two states will eventually come to the AVAILABLE (p. 108) state. You can use a storage volume during the PASS THROUGH and BOOTSTRAPPING states. However, you cannot take snapshots of the volume in these states.

### To remove an underlying disk allocated for the upload buffer (VMware ESXi)

1. In the vSphere client, right-click the name of your gateway VM and click **Edit Settings**.
2. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk allocated as upload buffer space, and then click **Remove**.

   Verify that the **Virtual Device Node** value in the **Virtual Machine Properties** dialog box has the same value that you noted previously. Doing this helps ensure that you remove the correct disk.

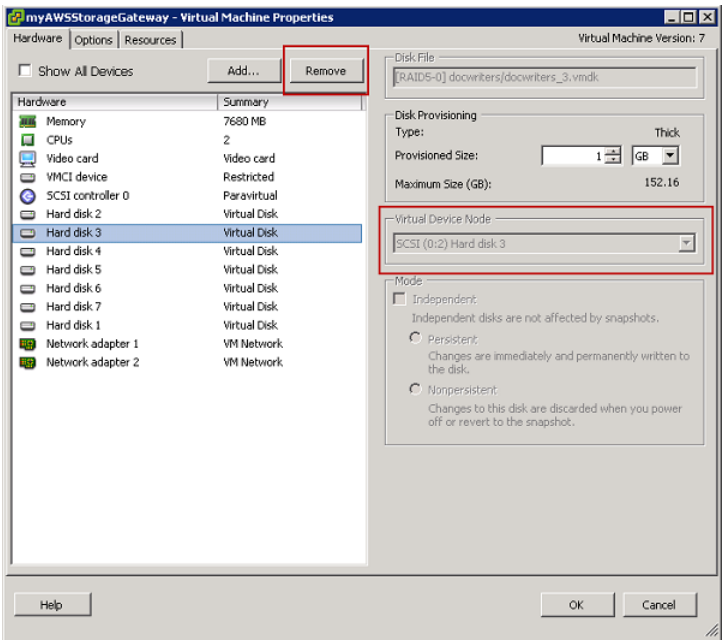


3. Choose an option in the **Removal Options** panel, and then click **OK** to complete the process of removing the disk.

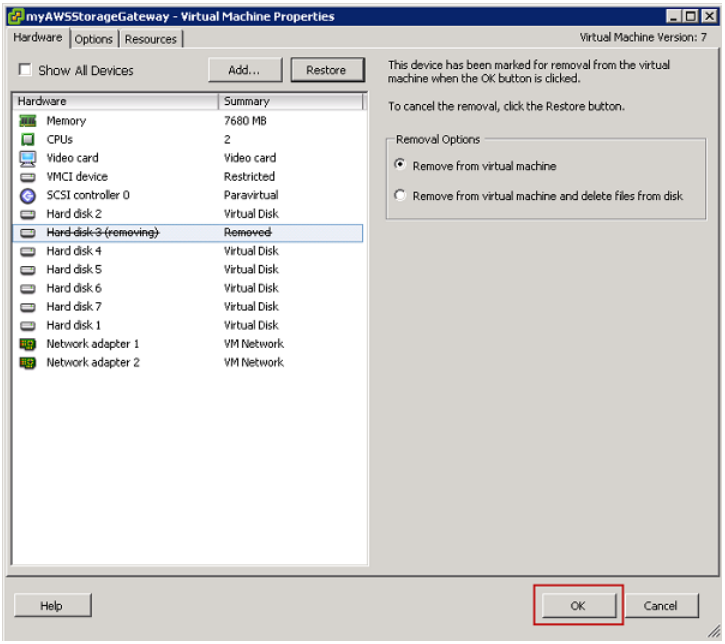


### To remove an underlying disk allocated for the upload buffer (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, right-click the name of your gateway VM and click **Settings**.
2. In the **Hardware** list of the **Settings** dialog box, select the disk to remove, and then click **Remove**.

   The disks you add to a gateway appear under the **SCSI Controller** entry in the **Hardware** list. Verify that the **Controller** and **Location** value are the same value that you noted previously. Doing this helps ensure that you remove the correct disk.

   The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.

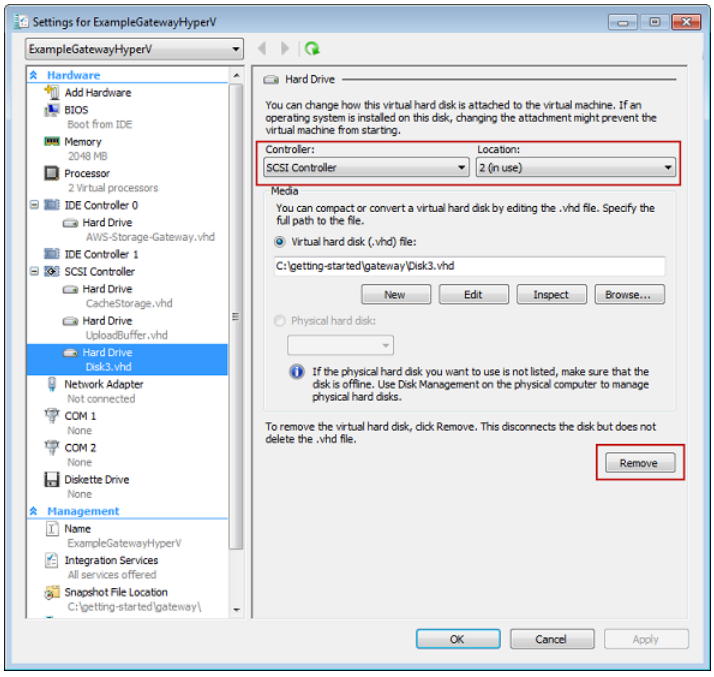

3. Click **OK** to apply the change.

## Adding Cache Storage (Gateway-Cached)

After you configure your initial gateway cache storage (see Configuring an Upload Buffer and Cache Storage (Gateway-Cached) (p. 177)), you can add cache storage to your gateway as your application needs change. To learn more about how to size your cache storage based on your application needs, see Managing Cache Storage (Gateway-Cached) (p. 172).

You can add more cache storage to your gateway without interrupting existing gateway functions and with the gateway VM turned on.

> **Note**
> Removing a disk allocated as cache storage is currently not supported.

You can add more cache storage by using the AWS Storage Gateway console or the AWS Storage Gateway API:

- For information on adding cache storage using the console, see To configure local disks as upload buffer and cache storage (p. 177). This procedure assumes that your activated gateway has at least one local disk available on its VM that you can allocate as cache storage for the gateway.
- For information on adding cache storage by using the API, see AddCache.

## Configuring an Upload Buffer and Cache Storage (Gateway-Cached)

As part of setting up your gateway-cached volume, you allocate storage (that is, disks) for your gateway. You use the suggested sizing formula in deciding the disk sizes. For more information on sizing the upload buffer, see Managing the Upload Buffer and Cache Storage (Gateway-Cached) (p. 169). The upload buffer and cache storage are created from local disks you provisioned for your gateway VM when you first created your gateway. After your gateway is up and running, you might decide to configure additional upload buffer or cache storage to your gateway. If you are configuring local storage for the first time, see Configure Local Storage and Alarms (Gateway-Cached Volumes) (p. 68) for instructions. The following procedure shows you how to configure an upload buffer and cached storage to your gateway.

### To configure local disks as upload buffer and cache storage

1. In the AWS Storage Gateway console, click the **Gateway** tab.
2. Click the **Configure Local Storage** tab.

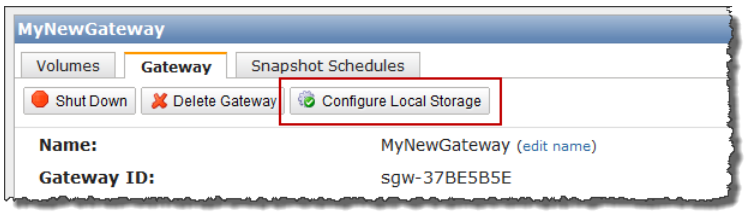


3. In the **Configure Your Activated Gateway** wizard, verify that there are local disks available to configure as an upload buffer and cache storage.

   The wizard shows a list of available disks on your local VM. If there are no local disks available, you must first add disks to your gateway-VTL VM. For more information on doing so, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23).

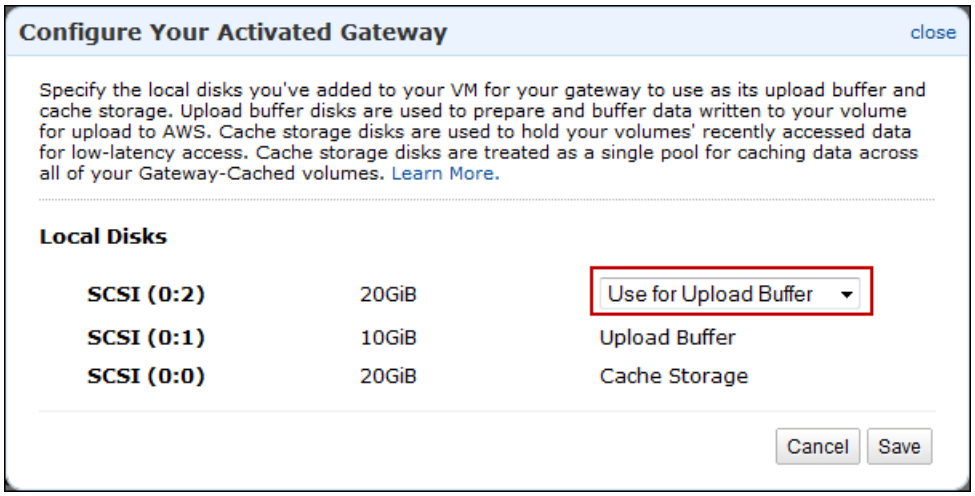
Configure Your Activated Gateway
close
Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage. Upload buffer disks are used to prepare and buffer data written to your volume for upload to AWS. Cache storage disks are used to hold your volumes' recently accessed data for low-latency access. Cache storage disks are treated as a single pool for caching data across all of your Gateway-Cached volumes. Learn More.
Local Disks
SCSI (0:2)
20GiB
Use for Upload Buffer
SCSI (0:1)
10GiB
Upload Buffer
SCSI (0:0)
20GiB
Cache Storage
Cancel
Save

**Important**
After configuring a disk as upload buffer storage, you lose any pre-existing data on the disk.

4. For each disk that you want to use, select **Use for Upload Buffer** or **Use for Cache Storage** from the drop-down list beside its name.
5. Click **Save** to save your configuration.

# Managing the Upload Buffer (Gateway-Stored)

As part of your initial deployment, you configure an upload buffer for your gateway. As you add volumes to provide more storage for your application data, you might need to add more upload buffer capacity to the gateway. You might also need to remove a local disk allocated as upload buffer—for example, because you want to replace a local disk that failed. This section shows how to determine if you need to add more upload buffer capacity and how to do it.

**Topics**


- Sizing the Upload Buffer (p. 179)
- Adding and Removing Upload Buffer Capacity (Gateway-Stored) (p. 180)
- Configuring an Upload Buffer (Gateway-Stored) (p. 184)


Your gateway uses the upload buffer to temporarily buffer your volume data prior to uploading it to AWS. The following diagram highlights the upload buffer in the larger picture of the AWS Storage Gateway gateway-stored architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 2).

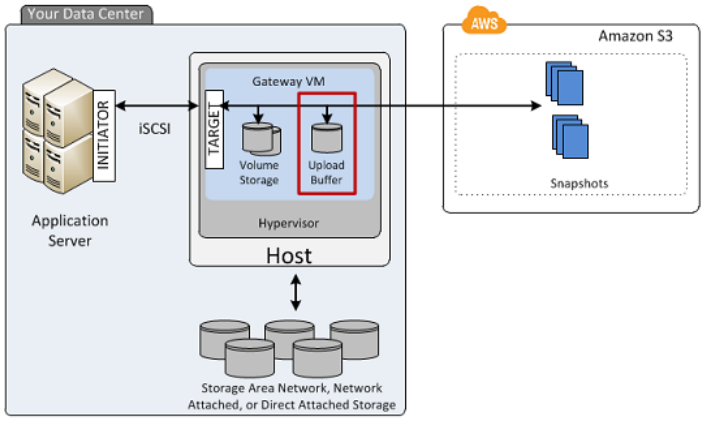


The amount of upload buffer space your gateway requires depends on several factors, such as the rate of incoming data to the storage volumes, the rate of outgoing data to AWS, and your network bandwidth. If your applications continue to write data at a fast rate to your storage volumes, and network throughput is not sufficient for the gateway to upload data to AWS, then eventually your upload buffer will be filled with data waiting to be uploaded to AWS. Here is some guidance you can follow to avoid this situation:

- **Use the sizing formula.** – As your application needs change, you should periodically review the recommended formula for sizing the upload buffer. For more information, see Sizing the Upload Buffer (p. 171).
- **Use Amazon CloudWatch metrics.** – You can proactively avoid the upload buffer from filling up by monitoring the percentage of upload buffer space your gateway is using in time. Amazon CloudWatch provides metrics such as the `UploadBufferPercentUsed` metric for monitoring your gateway's upload buffer usage (for more information, see Monitoring the Upload Buffer (p. 191)). You can set a threshold to trigger a notification to you when upload buffer usage exceeds a threshold. If your upload buffer is getting filled close to capacity, consider adding more buffer capacity to the gateway. For a full list of AWS Storage Gateway metrics, see Understanding AWS Storage Gateway Metrics (p. 195).

  You can see the current percentage of upload buffer usage on the **Gateway** tab of the AWS Storage Gateway console.
- **Monitor volume status.** – A volume status can indicate an issue with the upload buffer. For example, if the upload buffer capacity is reached, the storage volumes affected go into a PASS THROUGH (p. 108) status. Although your applications can continue to operate, writing and reading data to and from your storage volume, snapshots are not taken during this time.
- **Optimize your environment.** – If the speed of incoming write operations is too high compared to the outgoing network bandwidth, then the gateway might never catch up, no matter how much upload buffer capacity you provision. In this case, consider optimizing your gateway for better performance. For more information, see Optimizing Gateway Performance (p. 158).

## Sizing the Upload Buffer

You can determine the size of your upload buffer by using an upload buffer formula. We strongly recommend that you allocate at least 150 GiB of upload buffer. If the formula returns a value less than 150 GiB, use 150 GiB as the amount you allocate to the upload buffer. You can configure up to 2 TiB of upload buffer capacity per gateway.

> **Note**
> When the upload buffer reaches its capacity, your volume goes to PASS THROUGH status. In this status, new data that your application writes is persisted locally but not uploaded to AWS immediately. Thus, you cannot take new snapshots. When the upload buffer capacity frees up, the volume goes through BOOTSTRAPPING status. In this status, any new data that was persisted locally is uploaded to AWS. Finally, the volume returns to ACTIVE status. AWS Storage

Gateway then resumes normal synchronization of the data stored locally with the copy stored in AWS, and you can start taking new snapshots. For more information about volume status, see Understanding Storage Volume Status (p. 105).

To estimate the amount of upload buffer, you can determine the expected incoming and outgoing data rates and plug them into the following formula.

**Rate of incoming data**
This rate refers to the application throughput, the rate at which your on-premises applications write data to your gateway over some period of time.

**Rate of outgoing data**
This rate refers to the network throughput, the rate at which your gateway is able to upload data to AWS. This rate depends on your network speed, utilization, and whether you've enabled bandwidth throttling. This rate should be adjusted for compression. When uploading data to AWS, the gateway applies data compression where possible. For example, if your application data is text-only, you might get an effective compression ratio of about 2:1. However, if you are writing videos, the gateway might not be able to achieve any data compression and might require more upload buffer for the gateway.

If your incoming rate is higher than the outgoing rate, or if the formula returns a value less than 150 GiB, we strongly recommend that you allocate at least 150 GiB of upload buffer space.

For example, assume that your business applications will write text data to your gateway at a rate of 40 MB per second for 12 hours a day and your network throughput is 12 MB per second. Assuming a compression factor of 2:1 for the text data, you need to allocate approximately 690 GB of space for the upload buffer.

```
((40 MB/sec) - (12 MB/sec * 2)) * (12 hours * 3600 seconds/hour) = 691200
megabytes
```

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway as upload buffer space. Add more upload buffer space as needed using the AWS Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see Monitoring the Upload Buffer (p. 355).

If you decide that you need to change your upload buffer capacity, take one of the following actions.

| To... | Do This... |
|---|---|
| Add more upload buffer capacity to your gateway. | Follow the steps in Adding Upload Buffer Capacity (p. 181). |
| Remove a disk allocated as upload buffer. | Follow the steps in Removing Upload Buffer Capacity (Gateway-Stored) (p. 181). |

## Adding and Removing Upload Buffer Capacity (Gateway-Stored)

After you configure your initial gateway upload buffer (see Configuring an Upload Buffer (Gateway-Stored) (p. 184)), you can configure additional upload buffer capacity as your application needs

change. To learn more about how to size your upload buffer based on your application needs, see Sizing the Upload Buffer (p. 171).

You can add more buffer capacity to your gateway without interrupting existing gateway functions. Note that when you add more upload buffer capacity, you do so with the gateway VM turned on. However, when you reduce the amount of upload buffer capacity, you must first turn off the VM.

## Adding Upload Buffer Capacity

As your application needs change and you add more storage volume capacity, you might need to increase the gateway's upload buffer capacity. You can add more buffer capacity using the AWS Storage Gateway console or the AWS Storage Gateway API. For information on adding buffer capacity with the API, see AddUploadBuffer. The following procedure shows how to add more buffer capacity using the console. It assumes that your activated gateway has at least one local disk available on its VM that you can allocate as an upload buffer to the gateway.

### To configure additional upload buffer space for your gateway

- Follow the steps in To configure a local disk as an upload buffer for your gateway (p. 184).

## Removing Upload Buffer Capacity (Gateway-Stored)

As your application needs change and you change the storage volume configuration for a gateway, you might need to decrease the gateway's upload buffer capacity. Or, a local disk allocated to upload buffer might fail and you might need to remove that disk from the upload buffer and assign a new local disk. In both cases, you can remove buffer capacity using the AWS Storage Gateway console. The following procedure assumes that your activated gateway has at least one local disk allocated as an upload buffer to the gateway. In the procedure, you start on the AWS Storage Gateway console, leave the console and use the VMware vSphere client or the Microsoft Hyper-V Manager to remove the disk, and then return to the console.

### To find the ID of a disk allocated as an upload buffer

1. On the **Gateway** tab of the AWS Storage Gateway console, click **Configure Local Storage**.

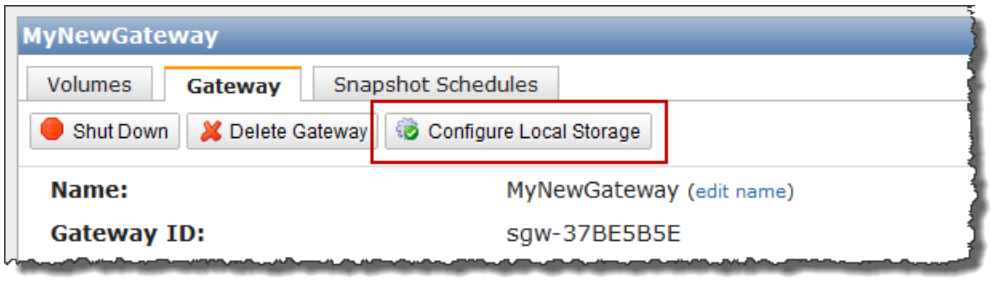


2. In the **Configure Local Upload Buffer** dialog box, note the value of the virtual device node for the local disk to be removed. You can find the node value in the **Local Disks** column. For example, in the following dialog box, the device node **SCSI (0:2)** is highlighted.

   You use the disk's virtual device node in the hypervisor client to help ensure that you remove the correct disk.

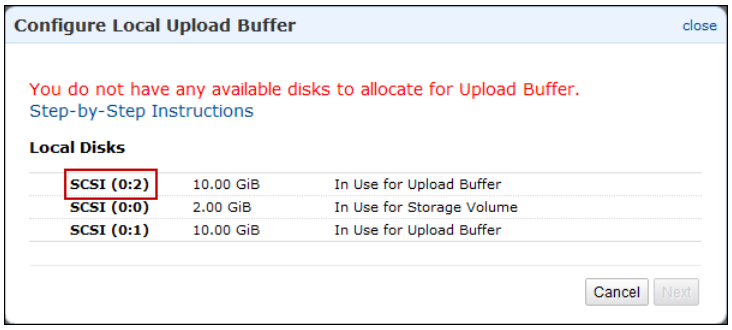
Configure Local Upload Buffer
close
You do not have any available disks to allocate for Upload Buffer.
Step-by-Step Instructions
Local Disks
SCSI (0:2) 10.00 GiB In Use for Upload Buffer
SCSI (0:0) 2.00 GiB In Use for Storage Volume
SCSI (0:1) 10.00 GiB In Use for Upload Buffer
Cancel
Next

3. Shut down the gateway by following the steps in the Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console (p. 127) procedure.

   **Note**
   Before shutting down the gateway, ensure that no application is writing data to it and that no snapshots are in progress. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the AWS Storage Gateway console. For more information, see Editing a Snapshot Schedule (p. 210).

4. To remove the underlying local disk, do one of the following procedures.

| For a Gateway Hosted In... | Do This... |
|---|---|
| VMware ESXi | Follow the steps in To remove the underlying disk allocated as an upload buffer (VMware ESXi) (p. 182). |
| Microsoft Hyper-V | Follow the steps in To remove the underlying disk allocated as an upload buffer (Microsoft Hyper-V) (p. 183). |

5. On the AWS Storage Gateway console, turn on the gateway.

   **Important**
   After removing a disk used as an upload buffer, you must turn the gateway back on before adding new disks to the VM.

6. In the **Volumes** tab of the AWS Storage Gateway console, check that all storage volumes have a status of AVAILABLE (p. 108).

   After a gateway restart, a storage volume might go through the PASS THROUGH (p. 108) and BOOTSTRAPPING (p. 108) states as the gateway adjusts to the upload buffer disk that you removed. A storage volume that passes through these two states will eventually come to the AVAILABLE (p. 108) state. You can use a storage volume during the pass through and bootstrapping states. However, you cannot take snapshots of the volume in these states.

### To remove the underlying disk allocated as an upload buffer (VMware ESXi)

1. In the vSphere client, right-click the name of your gateway VM and click **Edit Settings**.
2. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk allocated as an upload buffer, and then click **Remove**.

   Verify that the **Virtual Device Node** value in the **Virtual Machine Properties** dialog box has the same value that you noted previously. Doing this helps ensure that you remove the correct disk.

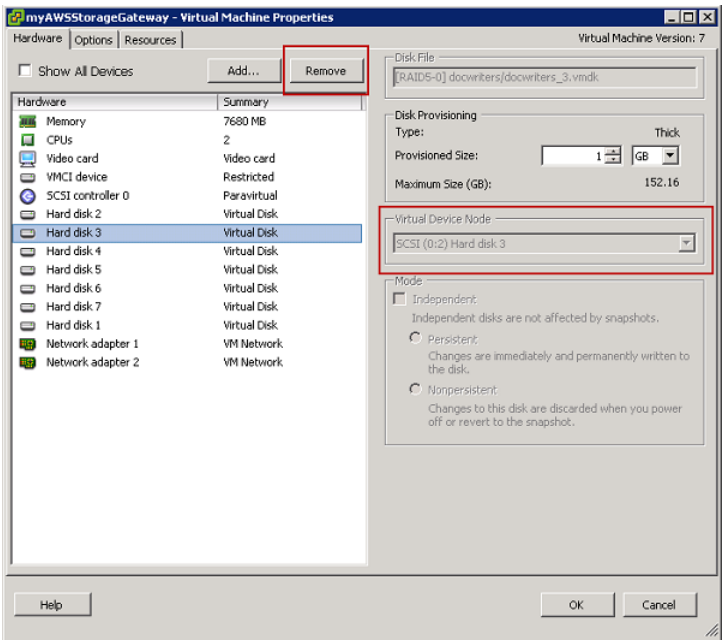


3. Choose an option in the **Removal Options** panel, and then click **OK** to complete the process of removing the disk.

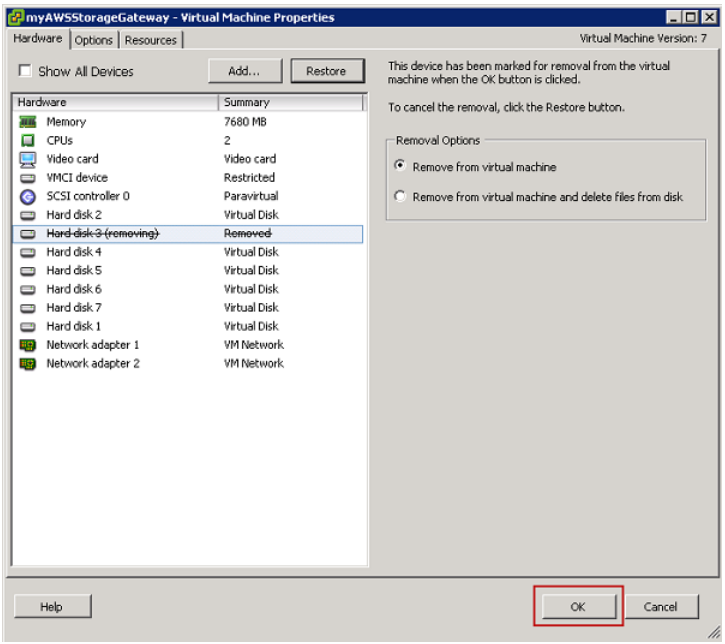


### To remove the underlying disk allocated as an upload buffer (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, right-click the name of your gateway VM and click **Settings**.
2. In the **Hardware** list of the **Settings** dialog box, select the disk to remove, and then click **Remove**.

   The disks you add to a gateway appear under **SCSI Controller** in the **Hardware** list. Verify that the **Controller** and **Location** values correspond to the virtual device node value that you noted previously. Doing this helps ensure that you remove the correct disk.

The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.

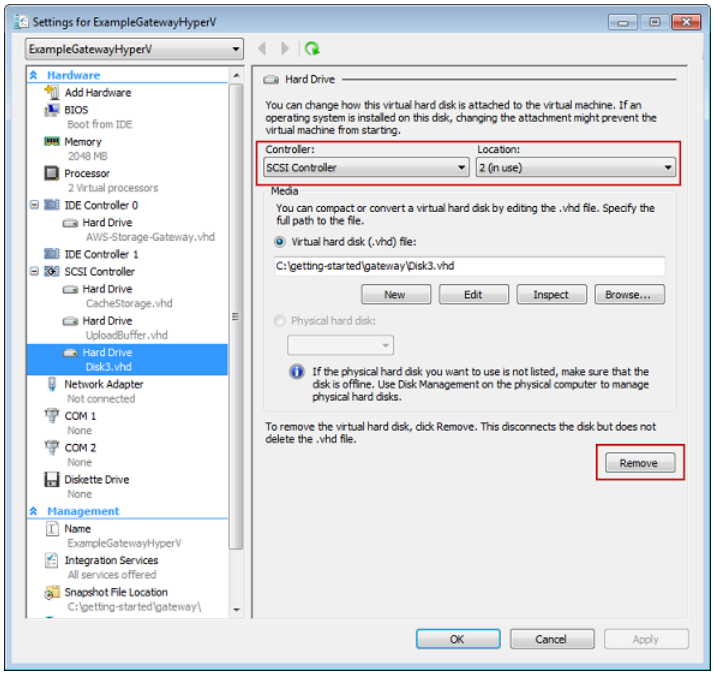


3. Click **OK** to apply the change.

## Configuring an Upload Buffer (Gateway-Stored)

As part of setting up your gateway-stored volumes, you allocate storage (that is, disks) for your gateway. You use the suggested sizing formula in deciding the disk sizes. For more information on sizing the upload buffer, see Sizing Cache Storage (p. 173). The upload buffer and cache storage are created from local disks you provisioned for your gateway VM when you first created your gateway. After your gateway is up and running, you might decide to configure additional upload buffer or cache storage for your gateway. If you are configuring local storage for the first time, see Configure Local Storage and Alarms (Gateway-Stored Volumes) (p. 71) for instructions. The following procedure shows you how to configure an upload buffer for your gateway.

### To configure a local disk as an upload buffer for your gateway

1. In the AWS Storage Gateway console, click the **Gateway** tab.
2. Click the **Configure Local Storage** tab.

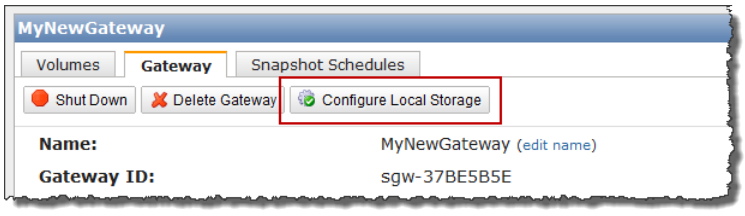


3. In the **Configure Your Activated Gateway** wizard, verify that there are local disks available to configure as an upload buffer and cache storage.

   The wizard shows a list of available disks on your local VM. If there are no local disks available, you must first add disks to your gateway VM. For more information on doing so, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23).

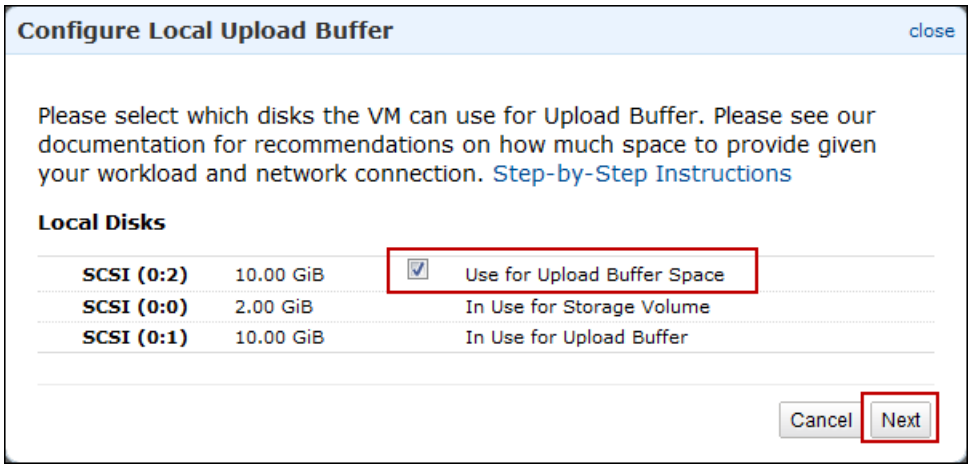
Configure Local Upload Buffer
close
Please select which disks the VM can use for Upload Buffer. Please see our documentation for recommendations on how much space to provide given your workload and network connection. Step-by-Step Instructions
Local Disks
SCSI (0:2) 10.00 GiB Use for Upload Buffer Space
SCSI (0:0) 2.00 GiB In Use for Storage Volume
SCSI (0:1) 10.00 GiB In Use for Upload Buffer
Cancel
Next

**Important**
After configuring a disk as upload buffer storage, you lose any pre-existing data on the disk.

4. Click to select the check box next to the disks that you want to allocate to the gateway as the upload buffer, and then click **Next**. The **Next** button is enabled only if you select at least one disk.
5. In the confirmation dialog box, read the text, click to select the confirmation check box, and then click **Confirm**.

   Doing this allocates the disk as upload buffer for the gateway.

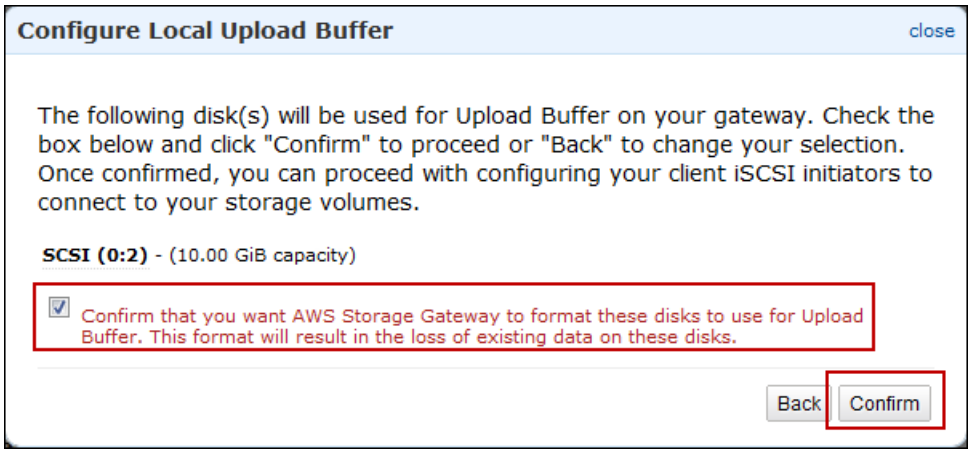
Configure Local Upload Buffer
close
The following disk(s) will be used for Upload Buffer on your gateway. Check the box below and click "Confirm" to proceed or "Back" to change your selection. Once confirmed, you can proceed with configuring your client iSCSI initiators to connect to your storage volumes.
SCSI (0:2) - (10.00 GiB capacity)
Confirm that you want AWS Storage Gateway to format these disks to use for Upload Buffer. This format will result in the loss of existing data on these disks.
Back
Confirm

# Monitoring Your Gateway

This section discusses how to monitor a gateway in a gateway-cached or gateway-stored setup, including monitoring the volumes associated with the gateway and the upload buffer. You use the AWS Management Console to view metrics for your gateway. For example, you can view the number of bytes used in read and write operations, the time spent in read and write operations, and the time to retrieve data from the AWS cloud. With metrics, you can track the health of your gateway and set up alarms to notify you when one or more metrics are outside a defined threshold.

**Topics**


- Using Amazon CloudWatch Metrics (p. 186)
- Measuring Performance Between Your Application and Gateway (p. 187)
- Measuring Performance Between Your Gateway and AWS (p. 188)

- Monitoring the Upload Buffer (p. 191)
- Monitoring Cache Storage (p. 194)
- Understanding AWS Storage Gateway Metrics (p. 195)

AWS Storage Gateway provides Amazon CloudWatch metrics at no additional charge. AWS Storage Gateway metrics are recorded for a period of two weeks. This functionality lets you access historical information and provides you with a better perspective of how your gateway and volumes are performing. For detailed information about Amazon CloudWatch, go to the Amazon CloudWatch Developer Guide.

# Using Amazon CloudWatch Metrics

You can get monitoring data for your gateway using either the AWS Management Console or the Amazon CloudWatch API. The console displays a series of graphs based on the raw data from the Amazon CloudWatch API. The Amazon CloudWatch API can be also be used, through one of the Amazon AWS Software Development Kits (SDKs) or the Amazon CloudWatch API tools. Depending on your needs, you might prefer to use either the graphs displayed in the console or retrieved from the API.

Regardless of which method you choose to use to work with metrics, you must specify the following information.

- First, you specify the metric dimension to work with. A *dimension* is a name-value pair that helps you to uniquely identify a metric. The dimensions for AWS Storage Gateway are `GatewayId`, `GatewayName`, and `VolumeId`. In the Amazon CloudWatch console, the `Gateway Metrics` and `Volume Metrics` views are provided to easily select gateway-specific and volume-specific dimensions. For more information about dimensions, see Dimensions in the *Amazon CloudWatch Developer Guide*.
- Second, you specify the metric name, such as `ReadBytes`.

The following table summarizes the types of AWS Storage Gateway metric data that are available to you.

| Amazon Cloud-Watch Namespace | Dimension | Description |
|---|---|---|
| `AWS/StorageGateway` | `GatewayId`, `GatewayName` | These dimensions filter for metric data that describes aspects of the gateway. You can identify a gateway to work with by specifying both the `GatewayId` and the `GatewayName` dimensions.<br><br>Throughput and latency data of a gateway is based on all the volumes in the gateway.<br><br>Data is available automatically in 5-minute periods at no charge. |
| | `VolumeId` | This dimension filters for metric data that is specific to a storage volume. Identify a storage volume to work with by its `VolumeId` dimension.<br><br>Data is available automatically in 5-minute periods at no charge. |

Working with gateway and volume metrics is similar to working with other service metrics. The Amazon CloudWatch documentation outlines many of the common tasks, listed following for your convenience:

- Listing Available Metrics

- Getting Statistics for a Metric
- Creating CloudWatch Alarms

# Measuring Performance Between Your Application and Gateway

Data throughput, data latency, and operations per second are three measures that you can use to understand how your application storage using your gateway is performing. When you use the correct aggregation statistic, these values can be measured by using the AWS Storage Gateway metrics that are provided for you.

A *statistic* is an aggregation of a metric over a specified period of time. When you view the values of a metric in Amazon CloudWatch, use the `Average` statistic for data latency (milliseconds), use the `Sum` statistic for data throughput (bytes per second), and use the `Samples` statistic for input/output operations per second (IOPS). For more information, see Statistics in the *Amazon CloudWatch Developer Guide*.

The following table summarizes the metrics and corresponding statistic you can use to measure the throughput, latency, and IOPS between your applications and gateways.

| Item of Interest | How to Measure |
|---|---|
| Throughput | Use the `ReadBytes` and `WriteBytes` metrics with the `Sum` Amazon CloudWatch statistic. For example, the `Sum` value of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput as a rate in bytes per second. |
| Latency | Use the `ReadTime` and `WriteTime` metrics with the `Average` Amazon CloudWatch statistic. For example, the `Average` value of the `ReadTime` metric gives you the latency per operation over the sample period of time. |
| IOPS | Use the `ReadBytes` and `WriteBytes` metrics with the `Samples` Amazon CloudWatch statistic. For example, the `Samples` value of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you IOPS. |

For the average latency graphs and average size graphs, the average is calculated over the total number of operations (read or write, whichever is applicable to the graph) that completed during the period.

The following procedures assume that you start in the Amazon CloudWatch console.

### To measure the data throughput from an application to a storage volume

1. Select the **StorageGateway: Volume Metrics** dimension, and find the storage volume that you want to work with.
2. Select the `ReadBytes` and `WriteBytes` metrics.
3. Select a **Time Range** value.
4. Select the `Sum` statistic.
5. Select a **Period** value of 5 minutes or greater.
6. In the resulting time-ordered sets of data points (one for `ReadBytes` and one for `WriteBytes`), divide each data point by the period (in seconds) to get the throughput at the sample point. The total throughput is the sum of the throughputs.

The following image shows the `ReadBytes` and `WriteBytes` metrics for a storage volume with the `Sum` statistic. In the image, the cursor over a data point displays information about the data point including its

value and the number of bytes. Divide the bytes value by the **Period** (5 minutes) to get the data throughput at that sample point. For the point highlighted, the read throughput is 2,384,199,680 bytes divided by 300 seconds, which is 7.6 megabytes per second.

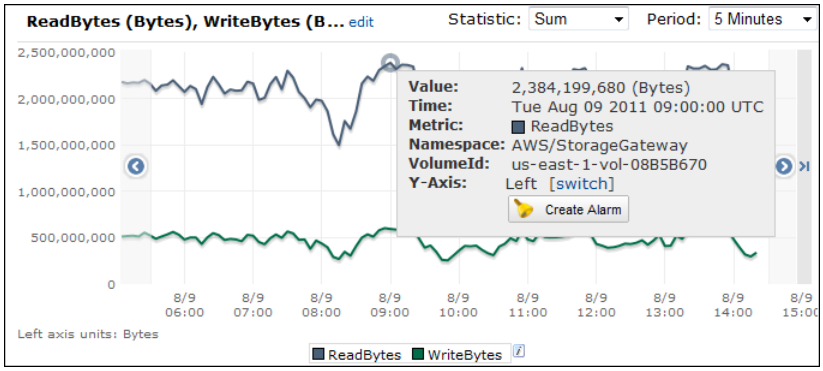


### To measure the data input/output operations per second from an application to a storage volume

1. Select the **StorageGateway: Volume Metrics** dimension, and find the storage volume that you want to work with.
2. Select the `ReadBytes` and `WriteBytes` metrics.
3. Select a **Time Range**.
4. Select the `Samples` statistic.
5. Select a **Period** of 5 minutes or greater.
6. In the resulting time-ordered sets of data points (one for `ReadBytes` and one for `WriteBytes`), divide each data point by the period (in seconds) to get IOPS.

The following image shows the `ReadBytes` and `WriteBytes` metrics for a storage volume with the `Samples` statistic. In the image, the cursor over a data point displays information about the data point, including its value and the number of samples. Divide the samples value by the **Period** (5 minutes) to get the operations per second at that sample point. For the point highlighted, the number of write operations is 24,373 bytes divided by 300 seconds, which is 81 write operations per second.

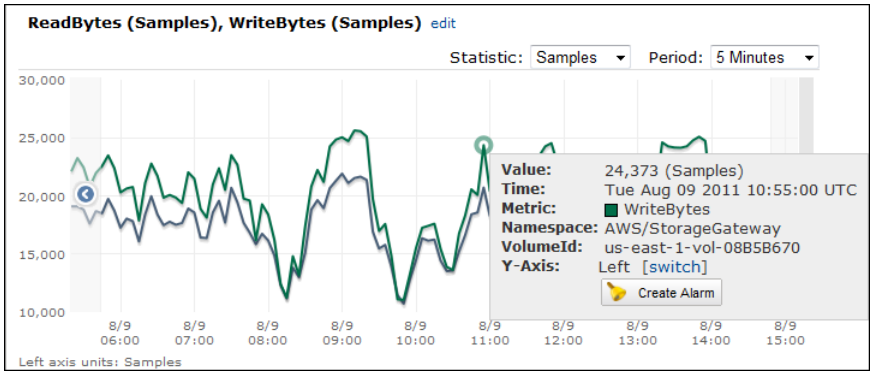


## Measuring Performance Between Your Gateway and AWS

Data throughput, data latency, and operations per second are three measures that you can use to understand how your application storage using the AWS Storage Gateway is performing. These three values can be measured using the AWS Storage Gateway metrics provided for you when you use the correct aggregation statistic. The following table summarizes the metrics and corresponding statistic to

use to measure the throughput, latency, and input/output operations per second (IOPS) between your gateway and AWS.

| Item of Interest | How to Measure |
| --- | --- |
| Throughput | Use the `ReadBytes` and `WriteBytes` metrics with the `Sum` Amazon CloudWatch statistic. For example, the `Sum` value of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput as a rate in bytes per second. |
| Latency | Use the `ReadTime` and `WriteTime` metrics with the `Average` Amazon CloudWatch statistic. For example, the `Average` value of the `ReadTime` metric gives you the latency per operation over the sample period of time. |
| IOPS | Use the `ReadBytes` and `WriteBytes` metrics with the `Samples` Amazon CloudWatch statistic. For example, the `Samples` value of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you IOPS. |
| Throughput to AWS | Use the `CloudBytesDownloaded` and `CloudBytesUploaded` metrics with the `Sum` Amazon CloudWatch statistic. For example, the `Sum` value of the `CloudBytesDownloaded` metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput from AWS to the gateway as bytes per second. |
| Latency of data to AWS | Use the `CloudDownloadLatency` metric with the `Average` statistic. For example, the `Average` statistic of the `CloudDownloadLatency` metric gives you the latency per operation. |

The following procedures assume that you start in the Amazon CloudWatch console.

### To measure the upload data throughput from a gateway to AWS

1. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.
2. Select the `CloudBytesUploaded` metric.
3. Select a **Time Range** value.
4. Select the `Sum` statistic.
5. Select a **Period** value of 5 minutes or greater.
6. In the resulting time-ordered set of data points, divide each data point by the period (in seconds) to get the throughput at that sample period.

The following image shows the `CloudBytesUploaded` metric for a gateway volume with the `Sum` statistic. In the image, the cursor over a data point displays information about the data point, including its value and bytes uploaded. Divide this value by the **Period** (5 minutes) to get the throughput at that sample point. For the point highlighted, the throughput from the gateway to AWS is 555,544,576 bytes divided by 300 seconds, which is 1.7 megabytes per second.

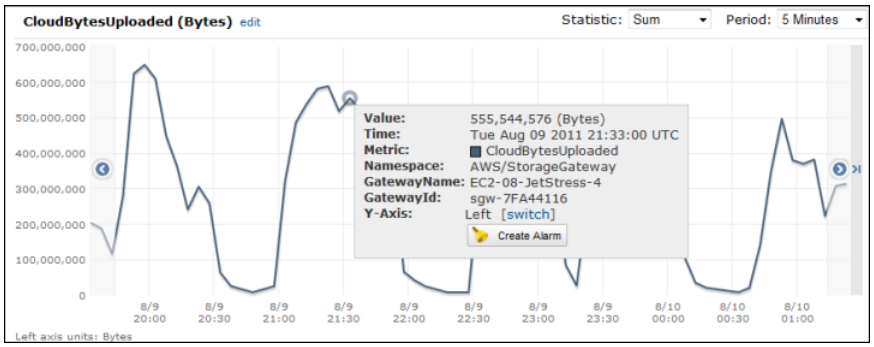


### To measure the latency per operation of a gateway

1. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.
2. Select the `ReadTime` and `WriteTime` metrics.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.
6. In the resulting time-ordered set of points (one for `ReadTime` and one for `WriteTime`), add data points at the same time sample to get to the total latency in milliseconds.

### To measure the data latency from a gateway to AWS

1. Select the **StorageGateway: GatewayMetrics** dimension, and find the gateway that you want to work with.
2. Select the `CloudDownloadLatency` metric.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the latency in milliseconds.

### To set an upper threshold alarm for a gateway's throughput to AWS

1. Start the Create Alarm Wizard.
2. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.
3. Select the `CloudBytesUploaded` metric.
4. To define the alarm, define the alarm state when the `CloudBytesUploaded` metric is greater than or equal to a specified value for a specified time. For example, you can define an alarm state when the `CloudBytesUploaded` metric is greater than 10 MB for 60 minutes.
5. Configure the actions to take for the alarm state. For example, you can have an email notification sent to you.
6. Create the alarm.

### To set an upper threshold alarm for reading data from AWS

1. Start the Create Alarm Wizard.

2. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.
3. Select the `CloudDownloadLatency` metric.
4. Define the alarm by defining the alarm state when the `CloudDownloadLatency` metric is greater than or equal to a specified value for a specified time. For example, you can define an alarm state when the `CloudDownloadLatency` is greater than 60,000 milliseconds for greater than 2 hours.
5. Configure the actions to take for the alarm state. For example, you can have an email notification sent to you.
6. Create the alarm. For information about how to create an alarm, see Step 2: Configure Local Storage and Alarms (Gateway-Cached and Gateway-Stored) (p. 67).

## Monitoring the Upload Buffer

The following section discusses how to monitor a gateway's upload buffer and how to create an alarm so that you get a notification when the buffer exceeds a specified threshold. This approach lets you proactively add buffer storage to a gateway before it fills completely and your storage application stops backing up to AWS.

Monitoring the upload buffer applies to both the gateway-cached and gateway-stored architectures. For more information, see How AWS Storage Gateway Works (Architecture) (p. 2)

> **Note**
> `WorkingStoragePercentUsed`, `WorkingStorageUsed`, and `WorkingStorageFree` metrics represent the upload buffer for only the gateway-stored volume setup prior to the release of the cached-volume feature in AWS Storage Gateway. Now you should use the equivalent upload buffer metrics `UploadBufferPercentUsed`, `UploadBufferUsed`, and `UploadBufferFree`. These metrics apply to both gateway architectures.

| Item of Interest | How to Measure |
|---|---|
| Upload buffer usage | Use the `UploadBufferPercentUsed`, `UploadBufferUsed`, and `UploadBufferFree` metrics with the `Average` statistic. For example, use the `UploadBufferUsed` with the `Average` statistic to analyze the storage usage over a time period. |

The following procedures assume that you start in the Amazon CloudWatch console.

### To measure upload buffer percent used

1. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.
2. Select the `UploadBufferPercentUsed` metric.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percent used of the upload buffer.

The following procedure shows you how to create an alarm using the Amazon CloudWatch console and the Create Alarm Wizard. To learn more about alarms and thresholds, see Creating CloudWatch Alarms.

### To set an upper threshold alarm for a gateway's upload buffer

1. Start the Create Alarm Wizard.

   a. In the Amazon CloudWatch console, click the **Alarms** link in the **Navigation** pane.
   b. In the **Your CloudWatch Alarms** pane, click **Create Alarm**.

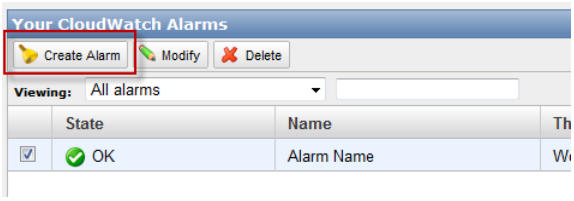


2. Specify a metric for your alarm.

   a. In the **SELECT METRIC** page of the Create Alarm Wizard, select the **AWS/StorageGateway:GatewayId,GatewayName** dimension, and then find the gateway that you want to work with.
   b. Select the `UploadBufferPercentUsed` metric. Use the `Average` statistic and a period of 5 minutes.

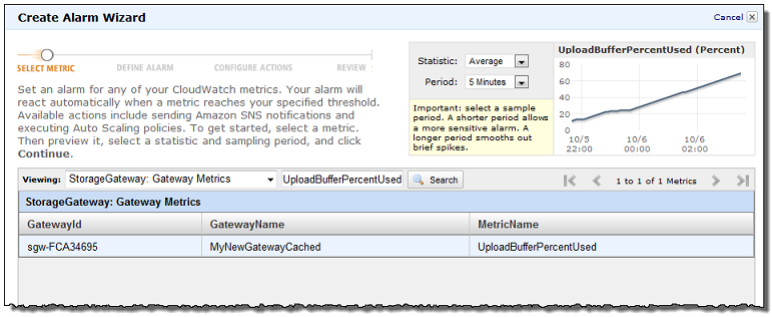
Create Alarm Wizard
Cancel
SELECT METRIC
DEFINE ALARM
CONFIGURE ACTIONS
REVIEW
Set an alarm for any of your CloudWatch metrics. Your alarm will react automatically when a metric reaches your specified threshold. Available actions include sending Amazon SNS notifications and executing Auto Scaling policies. To get started, select a metric. Then preview it, select a statistic and sampling period, and click Continue.
Statistic: Average
Period: 5 Minutes
Important: select a sample period. A shorter period allows a more sensitive alarm. A longer period smooths out brief spikes.
UploadBufferPercentUsed (Percent)
Viewing: StorageGateway: Gateway Metrics
UploadBufferPercentUsed
Search
1 to 1 of 1 Metrics
StorageGateway: Gateway Metrics
GatewayId
GatewayName
MetricName
sgw-FCA34695
MyNewGatewayCached
UploadBufferPercentUsed

   c. Click **Continue**.

3. Define the alarm name, description, and threshold.

   a. In the **DEFINE ALARM** page of the Create Alarm Wizard, identify your alarm by giving it a name and description in the **Name** and **Description** boxes.
   b. Define the alarm threshold.

      In the following image, the alarm state is defined for `UploadBufferPercentUsed` greater than or equal to 50 percent for 5 minutes.

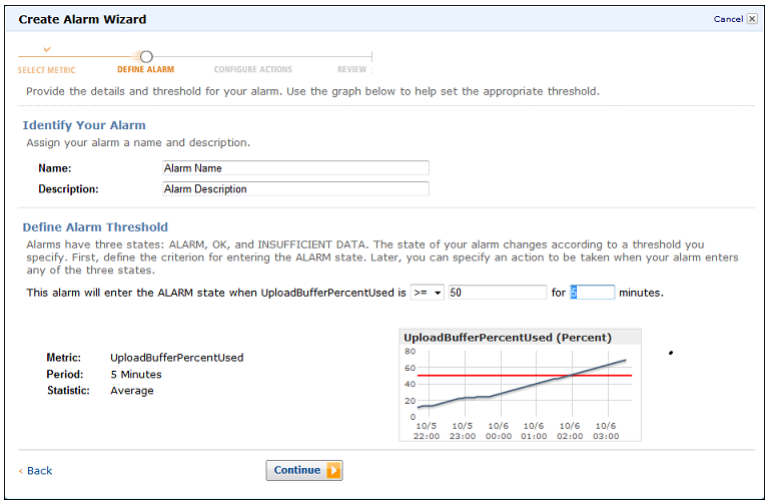
Create Alarm Wizard
Cancel
SELECT METRIC
DEFINE ALARM
CONFIGURE ACTIONS
REVIEW
Provide the details and threshold for your alarm. Use the graph below to help set the appropriate threshold.
Identify Your Alarm
Assign your alarm a name and description.
Name:
Alarm Name
Description:
Alarm Description
Define Alarm Threshold
Alarms have three states: ALARM, OK, and INSUFFICIENT DATA. The state of your alarm changes according to a threshold you specify. First, define the criterion for entering the ALARM state. Later, you can specify an action to be taken when your alarm enters any of the three states.
This alarm will enter the ALARM state when UploadBufferPercentUsed is >= 50 for 5 minutes.
Metric: UploadBufferPercentUsed
Period: 5 Minutes
Statistic: Average
UploadBufferPercentUsed (Percent)
80
60
40
20
0
10/5 22:00
10/5 23:00
10/6 00:00
10/6 01:00
10/6 02:00
10/6 03:00
Back
Continue

c. Click **Continue**.

4. Configure an email action for the alarm.

   a. In the **CONFIGURE ACTIONS** page of the Create Alarm Wizard, select **ALARM** from the **Alarm State** drop-down list.

   b. Select **Select or create email topic** from the **Topic** drop-down list.

      To create an email topic means that you set up an Amazon Simple Notification Service (Amazon SNS) topic. For more information about Amazon SNS, go to Set Up Amazon SNS.

   c. In the **Topic** box, enter a descriptive name for the topic.

   d. Click **ADD ACTION**.

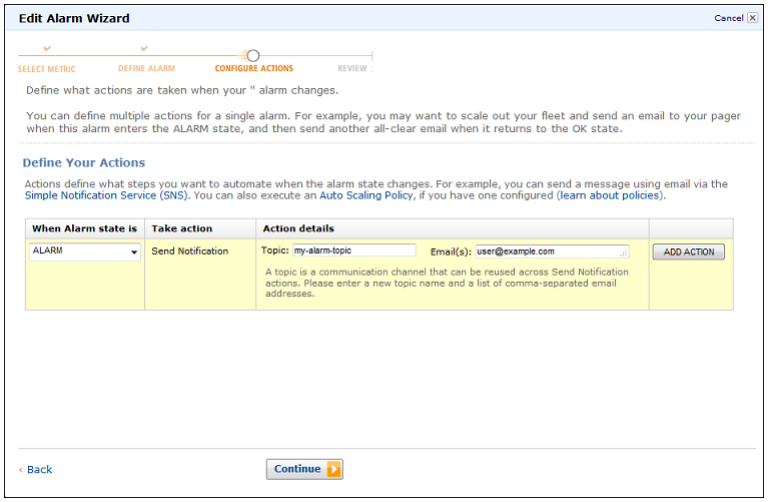
Edit Alarm Wizard
Cancel
SELECT METRIC
DEFINE ALARM
CONFIGURE ACTIONS
REVIEW
Define what actions are taken when your " alarm changes.
You can define multiple actions for a single alarm. For example, you may want to scale out your fleet and send an email to your pager when this alarm enters the ALARM state, and then send another all-clear email when it returns to the OK state.
Define Your Actions
Actions define what steps you want to automate when the alarm state changes. For example, you can send a message using email via the Simple Notification Service (SNS). You can also execute an Auto Scaling Policy, if you have one configured (learn about policies).
When Alarm state is
Take action
Action details
ALARM
Send Notification
Topic: my-alarm-topic
Email(s): user@example.com
ADD ACTION
A topic is a communication channel that can be reused across Send Notification actions. Please enter a new topic name and a list of comma-separated email addresses.
Back
Continue

e. Click **Continue**.

5. Review the alarm settings, and then create the alarm.

   a. In the **REVIEW** page of the Create Alarm Wizard, review the alarm definition, metric, and associated actions from this step. Associated actions include, for example, sending an email notification.

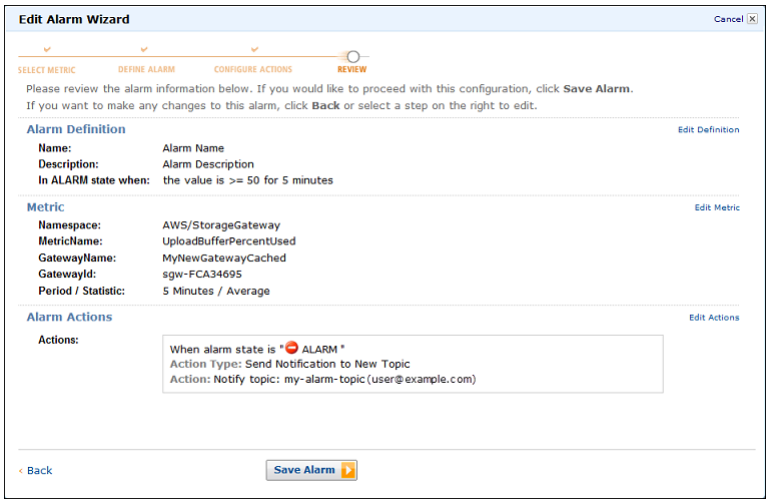
Edit Alarm Wizard
Cancel
SELECT METRIC
DEFINE ALARM
CONFIGURE ACTIONS
REVIEW
Please review the alarm information below. If you would like to proceed with this configuration, click Save Alarm.
If you want to make any changes to this alarm, click Back or select a step on the right to edit.
Alarm Definition
Edit Definition
Name: Alarm Name
Description: Alarm Description
In ALARM state when: the value is >= 50 for 5 minutes
Metric
Edit Metric
Namespace: AWS/StorageGateway
MetricName: UploadBufferPercentUsed
GatewayName: MyNewGatewayCached
GatewayId: sgw-FCA34695
Period / Statistic: 5 Minutes / Average
Alarm Actions
Edit Actions
Actions:
When alarm state is "ALARM"
Action Type: Send Notification to New Topic
Action: Notify topic: my-alarm-topic (user@example.com)
Back
Save Alarm

b. After reviewing the alarm summary, click **Save Alarm**.

6. Confirm your subscription to the alarm topic.

   a. Open the Amazon Simple Notification Service (Amazon SNS) email topic that is sent to the email address that you specified when creating the topic.

      The following image shows a notification.

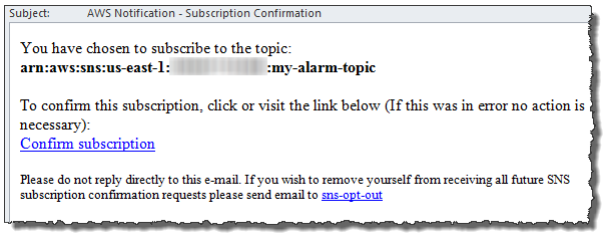
Subject: AWS Notification - Subscription Confirmation
You have chosen to subscribe to the topic:
arn:aws:sns:us-east-1:[redacted]:my-alarm-topic
To confirm this subscription, click or visit the link below (If this was in error no action is necessary):
Confirm subscription
Please do not reply directly to this e-mail. If you wish to remove yourself from receiving all future SNS subscription confirmation requests please send email to sns-opt-out

   b. Confirm your subscription by clicking the link in the email.

      A subscription confirmation appears.

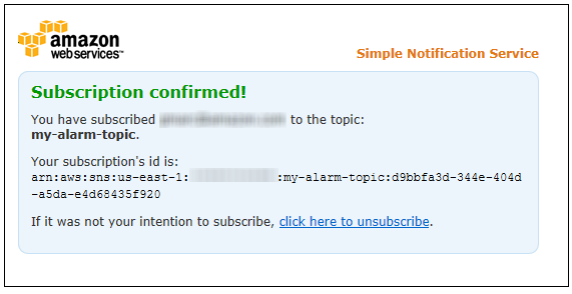
amazon web services
Simple Notification Service
Subscription confirmed!
You have subscribed [redacted] to the topic:
my-alarm-topic.
Your subscription's id is:
arn:aws:sns:us-east-1:[redacted]:my-alarm-topic:d9bbfa3d-344e-404d-a5da-e4d68435f920
If it was not your intention to subscribe, click here to unsubscribe.

## Monitoring Cache Storage

The following section discusses how to monitor a gateway's cache storage and how to create an alarm so that you get a notification when parameters of the cache pass specified thresholds. Using this alarm, you know when to proactively add cache storage to a gateway.

Monitoring cache storage applies only to the gateway-cached architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 2).

| Item of Interest | How to Measure |
|---|---|
| Total usage of cache | Use the `CachePercentUsed` and `TotalCacheSize` metrics with the `Average` statistic. For example, use the `CachePercentUsed` with the `Average` statistic to analyze the cache usage over a period of time.<br><br>The `TotalCacheSize` metric changes only when you add cache to the gateway. |
| Percentage of read requests that are served from the cache | Use the `CacheHitPercent` metric with the `Average` statistic.<br><br>Typically, you will want `CacheHitPercent` to remain high. |
| Percentage of cache that is dirty—that is, it contains content that has not been uploaded to AWS | Use the `CachePercentDirty` metrics with the `Average` statistic.<br><br>Typically, you want `CachePercentDirty` to remain low. |

The following procedures assume that you start in the Amazon CloudWatch console.

### To measure the cache's percentage dirty for a gateway and all its volumes

1. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.
2. Select the `CachePercentDirty` metric.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percentage of the cache that is dirty over the 5 minutes.

### To measure the cache's percentage dirty for a volume

1. Select the **StorageGateway: Volume Metrics** dimension, and find the volume that you want to work with.
2. Select the `CachePercentDirty` metric.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percentage of the cache that is dirty over the 5 minutes.

## Understanding AWS Storage Gateway Metrics

The following section discuses metrics you can use to monitor your gateway and volumes.

## Gateway Metrics

For the discussion in this topic, we define *gateway* metrics as metrics that are scoped to the gateway—that is, they measure something about the gateway. Because a gateway contains one or more volumes, a gateway-specific metric is representative of all volumes on the gateway. For example, the `CloudBytesUploaded` metric is the total number of bytes that the gateway sent to the cloud during the reporting period. This metric includes the activity of all the volumes on the gateway.

When working with gateway metric data, you will specify the unique identification of the gateway that you are interested in viewing metrics for. To do this, you specify both the `GatewayId` and the `GatewayName` values. When you want to work with metric for a gateway, you specify the gateway *dimension* in the metrics namespace, which distinguishes a gateway-specific metric from a volume-specific metric. For more information, see Using Amazon CloudWatch Metrics (p. 186).

**Topics**

The following table describes the AWS Storage Gateway metrics that you can use to get information about your gateway. The entries in the table are grouped functionally by measure.

> **Note**
> The reporting period for these metrics is 5 minutes.

| Metric | Description |
|---|---|
| `CacheHitPercent` | Percent of application reads served from the cache. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period.<br><br>Units: Percent |
| `CachePercentUsed` | Percent use of the gateway's cache storage. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period.<br><br>Units: Percent |
| `CachePercentDirty` | Percent of the gateway's cache that has not been persisted to AWS. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period.<br><br>Units: Percent |
| `CloudBytesDownloaded` | The total number of precompressed bytes that the gateway downloaded from AWS during the reporting period.<br><br>Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure input/output operations per second (IOPS).<br><br>Units: Bytes |
| `CloudDownloadLatency` | The total number of milliseconds spent reading data from AWS during the reporting period.<br><br>Use this metric with the `Average` statistic to measure latency.<br><br>Units: Milliseconds |

| Metric | Description |
| --- | --- |
| `CloudBytesUploaded` | The total number of precompressed bytes that the gateway uploaded to AWS during the reporting period.<br><br>Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS.<br><br>Units: Bytes |
| `UploadBufferFree` | The total amount of unused space in the gateway's upload buffer. The sample is taken at the end of the reporting period.<br><br>Units: Bytes |
| `UploadBufferPercentUsed` | Percent use of the gateway's upload buffer. The sample is taken at the end of the reporting period.<br><br>Units: Percent |
| `UploadBufferUsed` | The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period.<br><br>Units: Bytes |
| `QueuedWrites` | The number of bytes waiting to be written to AWS, sampled at the end of the reporting period for all volumes in the gateway. These bytes are kept in your gateway's working storage.<br><br>Units: Bytes |
| `ReadBytes` | The total number of bytes read from your on-premises applications in the reporting period for all volumes in the gateway.<br><br>Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS.<br><br>Units: Bytes |
| `ReadTime` | The total number of milliseconds spent to do read operations from your on-premises applications in the reporting period for all volumes in the gateway.<br><br>Use this metric with the `Average` statistic to measure latency.<br><br>Units: Milliseconds |
| `TotalCacheSize` | The total size of the cache in bytes. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period.<br><br>Units: Bytes |
| `WriteBytes` | The total number of bytes written to your on-premises applications in the reporting period for all volumes in the gateway.<br><br>Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS.<br><br>Units: Bytes |

| Metric | Description |
|---|---|
| `WriteTime` | The total number of milliseconds spent to do write operations from your on-premises applications in the reporting period for all volumes in the gateway.<br><br>Use this metric with the `Average` statistic to measure latency.<br><br>Units: Milliseconds |
| `WorkingStorageFree` | The total amount of unused space in the gateway's working storage. The sample is taken at the end of the reporting period.<br><br>**Note**<br>Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, `UploadBufferFree`.<br><br>Units: Bytes |
| `WorkingStoragePercentUsed` | Percent use of the gateway's upload buffer. The sample is taken at the end of the reporting period.<br><br>**Note**<br>Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, `UploadBufferPercentUsed`.<br><br>Units: Percent |
| `WorkingStorageUsed` | The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period.<br><br>**Note**<br>Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, `UploadBufferUsed`.<br><br>Units: Bytes |

## Storage Volume Metrics

In this section, we discuss the AWS Storage Gateway metrics that give you information about a storage volume of a gateway. Each volume of a gateway has a set of metrics associated with it. Note that some volume-specific metrics have the same name as certain gateway-specific metrics. These metrics represent the same kinds of measurements but are scoped to the volume instead of the gateway. You must always specify whether you want to work with either a gateway or a storage volume metric before working with a metric. Specifically, when working with volume metrics, you must specify the `VolumeId` that identifies the storage volume for which you are interested in viewing metrics. For more information, see Using Amazon CloudWatch Metrics (p. 186).

The following table describes the AWS Storage Gateway metrics that you can use to get information about your storage volumes.

| Metric | Description |
| --- | --- |
| `CacheHitPercent` | Percent of application read operations from the volume that are served from cache. This metric applies only to cached volumes. The sample is taken at the end of the reporting period.<br><br>When there are no application read operations from the volume, this metric reports 100 percent.<br><br>Units: Percent |
| `CachePercentUsed` | The volume's contribution to the overall percent use of the gateway's cache storage. This metric applies only to cached volumes. The sample is taken at the end of the reporting period.<br><br>Use the `CachePercentUsed` metric of the gateway to view overall percent use of the gateway's cache storage. For more information, see Gateway Metrics (p. 196).<br><br>Units: Percent |
| `CachePercentDirty` | The volume's contribution to the overall percentage of the gateway's cache that has not been persisted to AWS. This metric applies only to volumes in a gateway-cached setup. The sample is taken at the end of the reporting period.<br><br>Use the `CachePercentDirty` metric of the gateway to view the overall percentage of the gateway's cache that has not been persisted to AWS. For more information, see Gateway Metrics (p. 196).<br><br>Units: Percent |
| `ReadBytes` | The total number of bytes read from your on-premises applications in the reporting period.<br><br>Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS.<br><br>Units: Bytes |
| `ReadTime` | The total number of milliseconds spent to do read operations from your on-premises applications in the reporting period.<br><br>Use this metric with the `Average` statistic to measure latency.<br><br>Units: Milliseconds |
| `WriteBytes` | The total number of bytes written to your on-premises applications in the reporting period.<br><br>Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS.<br><br>Units: Bytes |
| `WriteTime` | The total number of milliseconds spent to do write operations from your on-premises applications in the reporting period.<br><br>Use this metric with the `Average` statistic to measure latency.<br><br>Units: Milliseconds |

| Metric | Description |
| --- | --- |
| `QueuedWrites` | The number of bytes waiting to be written to AWS, sampled at the end of the reporting period.<br><br>Units: Bytes |

# Logging Volume Gateway API Calls by Using AWS CloudTrail

AWS Storage Gateway is integrated with AWS CloudTrail, a service that captures API calls made by or on behalf of AWS Storage Gateway in your AWS account and delivers the log files to an Amazon S3 bucket that you specify. CloudTrail captures API calls from the AWS Storage Gateway console or from the AWS Storage Gateway API. Using the information collected by CloudTrail, you can determine what request was made to AWS Storage Gateway, the source IP address from which the request was made, who made the request, when it was made, and so on. To learn more about CloudTrail, including how to configure and enable it, see the *AWS CloudTrail User Guide*.

## AWS Storage Gateway Information in CloudTrail

When CloudTrail logging is enabled in your AWS account, API calls made to AWS Storage Gateway actions are tracked in log files. AWS Storage Gateway records are written together with other AWS service records in a log file. CloudTrail determines when to create and write to a new file based on a time period and file size.

All of the AWS Storage Gateway actions are logged and are documented in the Actions topic. For example, calls to the ActivateGateway, ListGateways, and ShutdownGateway actions generate entries in the CloudTrail log files.

Every log entry contains information about who generated the request. The user identity information in the log helps you determine whether the request was made with root or IAM user credentials, with temporary security credentials for a role or federated user, or by another AWS service. For more information, go to the **userIdentity** field in the CloudTrail Event Reference in the *AWS CloudTrail User Guide*.

You can store your log files in your bucket for as long as you want, but you can also define Amazon S3 lifecycle rules to archive or delete log files automatically. By default, your log files are encrypted by using Amazon S3 server-side encryption (SSE).

You can choose to have CloudTrail publish Amazon Simple Notification Service (Amazon SNS) notifications when new log files are delivered if you want to take quick action upon log file delivery. For more information, see Configuring Amazon SNS Notifications.

You can also aggregate AWS Storage Gateway log files from multiple AWS regions and multiple AWS accounts into a single Amazon S3 bucket. For more information, see Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket.

## Understanding AWS Storage Gateway Log File Entries

CloudTrail log files can contain one or more log entries where each entry is made up of multiple JSON-formatted events. A log entry represents a single request from any source and includes information about the requested action, any parameters, the date and time of the action, and so on. The log entries are not guaranteed to be in any particular order. That is, they are not an ordered stack trace of the public API calls.

The following example shows a CloudTrail log entry that demonstrates the ActivateGateway action.

```
      {"Records":[{"eventVersion":"1.02","userIdentity":{"type":"IAMUser","prin
cipalId":"AIDAII5AUEPBH2M7JTNVC","arn":"arn:aws:iam::111122223333:user/StorageG
ateway-team/JohnDoe",
       "accountId":"111122223333","accessKeyId":"AKIAIOSFODNN7EXAMPLE","user
Name":"JohnDoe"},"eventTime":"2014-12-04T16:19:00Z",
      "eventSource":"storagegateway.amazonaws.com","eventName":"ActivateGate
way","awsRegion":"us-east-1","sourceIPAddress":"192.0.2.0",
       "userAgent":"aws-cli/1.6.2 Python/2.7.6 Linux/2.6.18-164.el5",
       "requestParameters":{"gatewayTimezone":"GMT-5:00","gateway
Name":"cloudtrailgatewayvtl","gatewayRegion":"us-east-1","activationKey":"EHFBX-
1NDD0-P0IVU-PI259-DHK88","gatewayType":"VTL"},
       "responseElements":{"gatewayARN":"arn:aws:storagegateway:us-east-
1:111122223333:gateway/cloudtrailgatewayvtl"},"requestID":"54BT
FGNQI71987UJD2IHTCT8NF1Q8GLLE1QEU3KPGG6F0KSTAUU0",
       "eventID":"635f2ea2-7e42-45f0-bed1-8b17d7b74265","eventType":"AwsApiC
all","apiVersion":"20130630","recipientAccountId":"444455556666"}]}
```

The following example shows a CloudTrail log entry that demonstrates the ListGateways action.

```
      {"Records":[{"eventVersion":"1.02","userIdentity":{"type":"IAMUser","prin
cipalId":"AIDAII5AUEPBH2M7JTNVC","arn":"arn:aws:iam::111122223333:user/StorageG
ateway-team/JohnDoe",
       "accountId:"111122223333", "accessKeyId":"AKIAIOSFODNN7EXAMPLE","user
Name":"JohnDoe"},"eventTime":"2014-12-03T19:41:53Z",
       "eventSource":"storagegateway.amazonaws.com","eventName":"ListGate
ways","awsRegion":"us-east-1","sourceIPAddress":"192.0.2.0",
       "userAgent":"aws-cli/1.6.2 Python/2.7.6 Linux/2.6.18-164.el5",
       "requestParameters":null,"responseElements":null,"request
ID":"6U2N42CU37KAO8BG6V1I23FRSJ1Q8GLLE1QEU3KPGG6F0KSTAUU0",
       "eventID":"f76e5919-9362-48ff-a7c4-d203a189ec8d","eventType":"AwsApiC
all","apiVersion":"20130630",
       "recipientAccountId":"444455556666"}]}
```

# Working with Snapshots

AWS Storage Gateway provides the ability to back up point-in-time snapshots of your data to Amazon S3 for durable recovery. You can use the snapshot backups later on-premises or in Amazon Elastic Compute Cloud (Amazon EC2), and you can take snapshots on a one-time or scheduled basis. In this section, we show you the most common tasks that you can perform with snapshots, including creating a snapshot and restoring the snapshot to a volume that can then be mounted as an iSCSI device. When you restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume, the volume can then be attached to an Amazon EC2 instance.

AWS Storage Gateway continually and asynchronously uploads data to AWS to keep your local data synchronized with a copy stored in AWS. A benefit of this is that when snapshots are initiated, some or all the data has already been uploaded and snapshots complete quickly. Furthermore, snapshots are incremental—that is, the gateway uploads only the blocks of your volume that have changed since the last snapshot. For example, if you have 100 GiB of data and only 5 GiB data changed since the last

snapshot, then the gateway uploads only the 5 GiB of changed data. You can delete any snapshot. AWS Storage Gateway removes only the snapshot data that is not needed by other snapshots, letting you restore a volume from any of the active snapshots.

How snapshots can be effectively used in your AWS Storage Gateway setup depends on the type of gateway you set up—that is, gateway-cached or gateway-stored. For more information on gateway-cached and gateway-stored architecture, see How AWS Storage Gateway Works (Architecture) (p. 2)).

- For gateway-cached volumes, your volume data is already stored in Amazon S3. Snapshots can be used to preserve older versions of your data.
- For gateway-stored volumes, your volume data is stored on-premises. Snapshots provide durable, off-site backups in Amazon S3.

You can restore a snapshot to either a gateway volume or to an Amazon EBS volume. To restore a snapshot to a gateway storage volume, first set up and activate a local gateway. Getting Started with AWS Storage Gateway (p. 9) walks you through the steps of setting up a gateway. When you restore a snapshot to an Amazon EBS volume, you can then attach the EBS volume to an Amazon EC2 instance. For a complete list of charges and specific prices for Amazon EC2, go to the Amazon EC2 Pricing page.

**Topics**


- Summary of Snapshot Tasks (p. 202)
- Snapshot Consistency (p. 203)
- Finding a Snapshot (p. 203)
- Editing a Snapshot Schedule (p. 210)
- Creating an Ad Hoc Snapshot (p. 212)
- Deleting a Snapshot (p. 212)
- Restoring a Snapshot (p. 222)


# Summary of Snapshot Tasks

Because snapshots are key to using the AWS Storage Gateway service, you should understand at a high level what each snapshot operation does and why it does that. Each task is covered in detail in the linked section.

To work with the tasks, you should have one or more gateways that have been running for enough time so there are snapshots to work with. You can work with snapshots using the AWS Storage Gateway console, one of the AWS Software Development Kits (SDKs), or the AWS Storage Gateway REST API (see Operations in AWS Storage Gateway (p. 403)). In this section, we primarily show how to work with the console to perform gateway tasks.

| Snapshot Action | Common Scenarios |
|---|---|
| Finding | You might want to find a snapshot to see if it is complete, what time it was taken, what the size of the snapshot is, or the name of the volume the snapshot was taken from. For more information, see Finding a Snapshot (p. 203). |
| Scheduling | When you first set up a stored volume, a default snapshot schedule of once per day is set. You can change the frequency and timing of the snapshot schedule to fit your application needs. For more information, see Editing a Snapshot Schedule (p. 210). |

| Snapshot Action | Common Scenarios |
|---|---|
| Creating | Snapshots for stored volumes are automatically created on a schedule by default, and you can change the schedule of the snapshots as needed. You can also take an instantaneous snapshot at any time for both stored and cached volumes. For more information, see Creating an Ad Hoc Snapshot (p. 212). For information about troubleshooting snapshot issues, see Troubleshooting Snapshot Issues (p. 237). |
| Restoring | You can restore a snapshot to a new local volume, or you can use the snapshot to create an Amazon EBS volume and attach that to an Amazon EC2 instance. For more information, see Restoring a Snapshot to a Storage Volume (p. 222) and Restoring a Snapshot to an Amazon EBS Volume (p. 225). |
| Deleting | If you don't need a snapshot anymore, you can delete it. Because snapshots are incremental backups, if you delete a snapshot only the data that is not needed in other snapshots is deleted. For more information, see Deleting a Snapshot (p. 212). |

## Snapshot Consistency

Snapshots provide a point-in-time view of data that has been written to your AWS Storage Gateway volumes. However, snapshots only capture data that has been written to your storage volumes, which can exclude data that has been buffered by your client application or operating system. Your application and operating system will eventually flush this buffered data to your storage volumes. If you need to guarantee that your application data is flushed to disk prior to taking a snapshot, you should consult your specific application's documentation to understand if and how your application buffers data and how to flush this data.

If you need to guarantee that your operating system and file system have flushed their buffered data to disk prior to taking a snapshot, you can do this by taking your storage volume offline before taking a snapshot. Doing this forces your operating system to flush its data to disk. After the snapshot is complete, you can bring the volume back online. In Windows, use Disk Management (`diskmgmt.msc`) to select the storage volume and take it online or offline. To script this process in Windows, you can use a command line tool such as Diskpart.exe. In Linux, you can use the `mount` and `umount` commands.

## Finding a Snapshot

You will need to find a snapshot associated with a volume when you want to restore the snapshot to a new volume—for example, in a disaster recovery scenario—or to restore a previous version of your application data. To restore a snapshot, you need to know the snapshot ID. You can find a snapshot ID in several ways, including using the AWS Storage Gateway console, the Amazon EC2 console, or programmatically using one of the AWS Software Development Kits (SDKs).

If you list snapshots using the AWS Storage Gateway console, the list includes all snapshots generated from your gateway and also snapshots that you might have generated from Amazon EBS volumes. If you list snapshots on the Amazon EC2 console, more snapshot properties are shown to help you find your snapshot. The Amazon EC2 console also has a search filtering capability. Both console experiences are described following. In some scenarios, you might need to search using several snapshot properties at once—for example, status, start date, and description. In this case, you can use a programmatic approach. For examples, see Finding Snapshots Using the AWS SDK for Java (p. 205), Finding Snapshots Using the AWS SDK for .NET (p. 207), and Finding Snapshots Using the AWS Tools for Windows PowerShell (p. 209).

When you find your snapshot, you can view its details, including the date and time the snapshot was started and the storage volume on your gateway that was the source for the snapshot.

### To find a snapshot for a volume using the AWS Storage Gateway console

1. On the **AWS Storage Gateway** console, in the **Navigation** pane, click **Snapshots**.

   The **Snapshots** window shows a list of your snapshots.

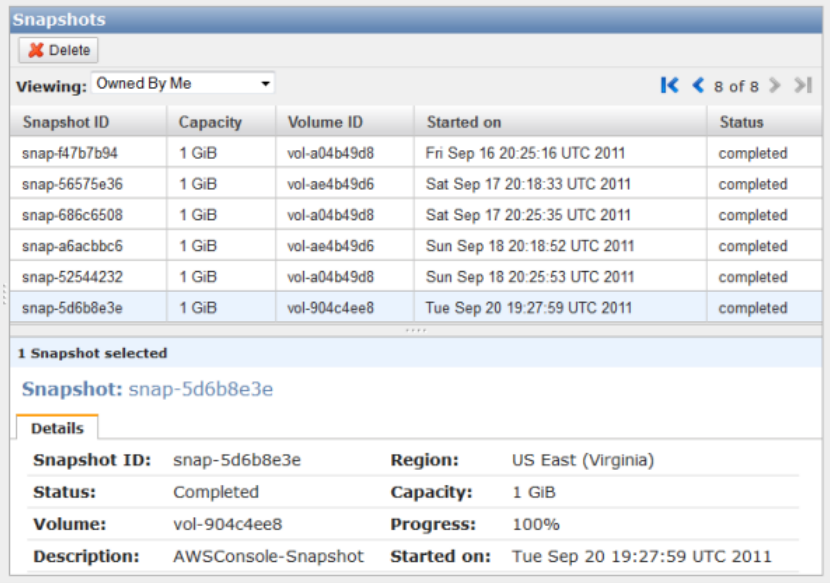


2. Find the snapshot that you are looking for by finding its volume ID in the **Volume ID** column.
3. Click the snapshot row to display the snapshot details.

### To find a snapshot for a volume using the Amazon EC2 Management console

1. On the Amazon EC2 Management console, in the **Navigation** pane, click **Snapshots** under **ELASTIC BLOCK STORE**.

   The **EBS Snapshots** window shows a list of your snapshots. This snapshot view offers more functionality for finding snapshots than on the AWS Storage Gateway console. In particular, this view shows descriptions that can be useful for filtering results—for instance, there is a pattern to snapshot descriptions that you can use to help find a snapshot. For example, in the console image following, note these points:

   - The row marked as 1 is a snapshot taken of an Amazon EBS volume. The name and description fields are specified in the snapshot creation. This snapshot is not from an AWS Storage Gateway operation.
   - The row marked as 2 is a one-time, ad hoc snapshot taken from the AWS Storage Gateway console. The descriptions for ad hoc snapshots contain "AWSConsole-Snapshot."
   - The row marked as 3 is a snapshot of a volume taken from a snapshot schedule by AWS Storage Gateway. The descriptions of scheduled snapshots contain the storage gateway ID, the volume ID, and the word "Schedule" in this pattern: *gatewayID*:*volumeID*:Schedule

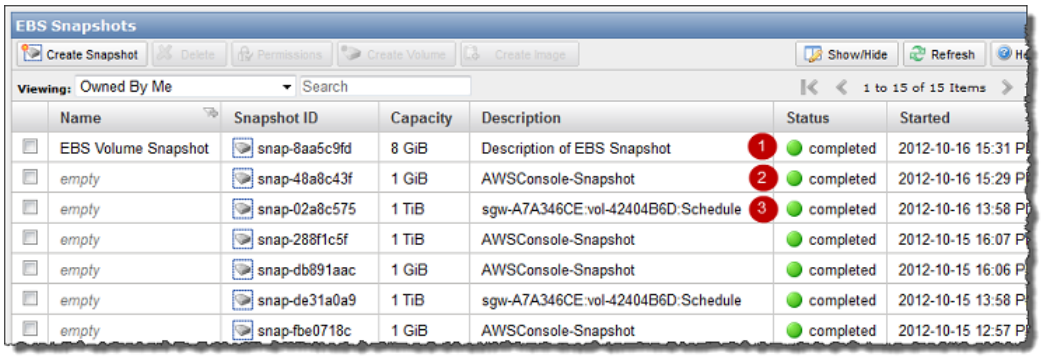


2. Find the snapshot that you are looking for in the list by typing some or all of the volume ID into the **Search** field.

   In the following example, only results containing **`vol-424`** are shown. You can find the volume ID on the AWS Storage Gateway console on the **Volumes** tab.

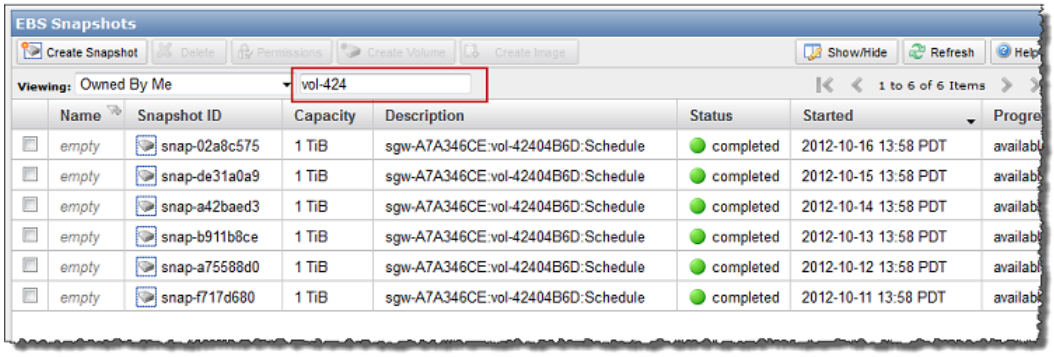


3. Click the snapshot row to display the snapshot details.

## Finding Snapshots Using the AWS SDK for Java

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS SDK for Java using several snapshot properties. To use the example code, you should be familiar with running a Java console application. For more information, go to Getting Started in the *AWS SDK for Java Developer Guide*.

### Example : Finding Snapshots Using the AWS SDK for Java

The following Java code example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses the AWS SDK for Java and the Amazon EC2 API. The Amazon EC2 API includes operations for working with snapshots.

You need to update the code and provide the service endpoint, a full or partial volume ID, a snapshot status, and a number of days to indicate a cutoff date—snapshots taken before the cutoff are not returned. For a list of AWS service endpoints you can use with Amazon EC2, go to Regions and Endpoints in the *Amazon Web Services Glossary*.

```
import java.io.IOException;
import java.util.Calendar;
import java.util.Date;
import java.util.GregorianCalendar;
import java.util.List;

import com.amazonaws.AmazonClientException;
import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.ec2.AmazonEC2Client;
import com.amazonaws.services.ec2.model.DescribeSnapshotsRequest;
import com.amazonaws.services.ec2.model.DescribeSnapshotsResult;
import com.amazonaws.services.ec2.model.Filter;
import com.amazonaws.services.ec2.model.Snapshot;


public class FindingSnapshotsExample {

    static AmazonEC2Client ec2Client;
    // A full volume id or partial fragment with "*".
    static String volumeID = "vol-424*";
    // Snapshot status to filter on: "completed", "pending", "error".
    static String status = "completed";
    // The number of days before which to not return snapshot results.
    static int daysBack = 10;
    // Service end point. Should be same region as volume/gateway.
   public static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";


    public static void main(String[] args) throws IOException {

        ec2Client = new AmazonEC2Client(new PropertiesCredentials(
                FindingSnapshotsExample.class.getResourceAsStream("AwsCreden
tials.properties")));
        ec2Client.setEndpoint(serviceURLEC2);

        FindingSnapshotsForAVolume();

    }

    private static void FindingSnapshotsForAVolume() {

        try {
            Filter[] filters = new Filter[2];
          filters[0] = new Filter().withName("volume-id").withValues(volumeID);

            filters[1] = new Filter().withName("status").withValues(status);
```

```
            DescribeSnapshotsRequest describeSnapshotsRequest =
                new DescribeSnapshotsRequest().withFilters(filters);
            DescribeSnapshotsResult describeSnapshotResult =
                ec2Client.describeSnapshots(describeSnapshotsRequest);

            List<Snapshot> snapshots = describeSnapshotResult.getSnapshots();
            System.out.println("volume-id = " + volumeID);
            for (Snapshot s : snapshots) {
                if (CompareDates(daysBack, s.getStartTime())) {
                    StringBuilder sb = new StringBuilder();
                    sb.append(s.getSnapshotId() + ", " + s.getStartTime() + ",
 " + s.getDescription());
                    System.out.println(sb.toString());
                }
            }
        } catch (AmazonClientException ace) {
            System.err.println(ace.getMessage());
        }
    }
    public static boolean CompareDates(int daysBack, Date snapshotDate) {
        Date today = new Date();
        Calendar cal = new GregorianCalendar();
        cal.setTime(today);
        cal.add(Calendar.DAY_OF_MONTH, -daysBack);
        Date cutoffDate = cal.getTime();
        return (snapshotDate.compareTo(cutoffDate) > 0) ? true : false;
    }

}
```

## Finding Snapshots Using the AWS SDK for .NET

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS SDK for .NET using several snapshot properties. To use the example code, you should be familiar with running a .NET console application. For more information, go to Getting Started in the *AWS SDK for .NET Developer Guide*.

### Example : Finding Snapshots Using the AWS SDK for .NET

The following C# code example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses the AWS SDK for .NET and the Amazon EC2 API. The Amazon EC2 API includes operations for working with snapshots.

> **Note**
> The following code example uses the AWS SDK for .NET version 1 and might include methods that are no longer available in the current version of the SDK. For more information, see Migrating Your Code to the Latest Version of the AWS SDK for .NET.

You need to update the code and provide the service endpoint, a full or partial volume ID, a snapshot status, and a number of days to indicate a cutoff date for the snapshots to return. For a list of AWS service endpoints you can use with Amazon EC2, go to Regions and Endpoints in the *Amazon Web Services Glossary*.

```
using System;
using System.Text;
using System.Collections.Generic;
using Amazon.EC2;
using Amazon.EC2.Model;

namespace AWSStorageGateway
{
    class FindingSnapshotsExample
    {
        static AmazonEC2Config ec2Config;
        static AmazonEC2Client ec2Client;
        // A full volume id or partial fragment with "*".
        static String volumeID = "vol-424*";
        // Snapshot status to filter on: "completed", "pending", "error".
        static String status = "completed";
        // The number of days before which to not return snapshot results.
        static int daysBack = 4;
        // Service endpoint. Should be same region as volume/gateway.
        static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";


        public static void Main(string[] args)
        {
            //Create a ec2 client
            ec2Config = new AmazonEC2Config();
            ec2Config.ServiceURL = serviceURLEC2;
            ec2Client = new AmazonEC2Client(ec2Config);

            FindingSnapshotsForAVolume();

            Console.WriteLine("\nTo continue, press Enter.");
            Console.Read();
        }

        private static void FindingSnapshotsForAVolume()
        {
            try
            {
                Filter[] filters = new Filter[2];
                filters[0] = new Filter().WithName("volume-id").With
Value(volumeID);
```

```
                filters[1] = new Filter().WithName("status").WithValue(status);


                 DescribeSnapshotsRequest describeSnapshotsRequest =
                     new DescribeSnapshotsRequest().WithFilter(filters);
                 DescribeSnapshotsResponse describeSnapshotsResponse =
                     ec2Client.DescribeSnapshots(describeSnapshotsRequest);

                 List<Snapshot> snapshots = describeSnapshotsResponse.Describe
SnapshotsResult.Snapshot;
                 Console.WriteLine("volume-id = " + volumeID);
                 foreach (Snapshot s in snapshots)
                 {
                     if (CompareDates(daysBack, s.StartTime))
                     {
                         StringBuilder sb = new StringBuilder();
                         sb.Append(s.SnapshotId + ", " + s.StartTime + ", " +
s.Description);
                         Console.WriteLine(sb.ToString());
                     }
                 }
             }
             catch (AmazonEC2Exception ex)
             {
                 Console.WriteLine(ex.Message);
             }

         }

         public static Boolean CompareDates(int daysBack, String d)
         {
             DateTime snapshotDate = DateTime.Parse(d);
             DateTime cutoffDate = DateTime.Now.Add(new TimeSpan(-daysBack, 0,
0, 0));
             return (DateTime.Compare(snapshotDate, cutoffDate) < 0) ? true :
false;
         }
     }
}
```

## Finding Snapshots Using the AWS Tools for Windows PowerShell

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS Tools for Windows PowerShell using several snapshot properties. To use the example code, you should be familiar with running a PowerShell script. For more information, go to Getting Started in the *AWS Tools for Windows PowerShell User Guide*.

### Example : Finding Snapshots Using the AWS Tools for Windows PowerShell

The following PowerShell script example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses AWS Tools for Windows PowerShell cmdlets for Amazon EC2. The Amazon EC2 cmdlets include operations for working with snapshots.

You need to update the script and provide a full or partial volume ID, a snapshot status, and a number of days to indicate a cutoff date for the snapshots to return.

```
<#
.DESCRIPTION
    Finds snapshots for a given volume and criteria about the snapshot.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from http://console.amazonaws.cn/powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
   For more info, see http://docs.amazonaws.cn/powershell/latest/userguide//spe
cifying-your-aws-credentials.html
.EXAMPLE
    powershell.exe .\SG_FindSnapshots.ps1
#>

# Criteria to use to filter the results returned.
$volumeID = "vol-424*"
$status = "completed"
$daysBack = 4

# Define filters.
$filter1 = New-Object Amazon.EC2.Model.Filter
$filter1.Name = "volume-id"
$filter1.Value.Add($volumeID)

$filter2 = New-Object Amazon.EC2.Model.Filter
$filter2.Name = "status"
$filter2.Value.Add($status)

$snapshots = get-EC2Snapshot -Filter $filter1, $filter2
$count = 0

foreach ($s in $snapshots)
 {
   $d = ([DateTime]::Now).AddDays(-$daysBack)
   if ([DateTime]::Compare($d, $s.StartTime) -gt 0)
   {
        # Meets criteria.
       $count +=1
       $sb = $s.SnapshotId + ", " + $s.StartTime + ", " + $s.Description
       Write-Output($sb)
   }
}
Write-Output ("Found " + $count + " snapshots that matched the criteria.")
```

## Editing a Snapshot Schedule

For gateway-stored volumes, AWS Storage Gateway creates a default snapshot schedule of once a day. This schedule helps ensure that your gateway can keep up with the rate of incoming write operations on

your local storage volumes. You can change the default snapshot schedule by specifying either the time the snapshot occurs each day or the frequency (every 1, 2, 4, 8, 12, or 24 hours), or both.

For gateway-cached volumes, AWS Storage Gateway does not create a default snapshot schedule. However, you can set up a snapshot schedule at any time if you need to. For gateway-cached volumes, because your data is stored in Amazon S3, snapshots (and a snapshot schedule) are not needed for disaster recovery purposes.

In the following steps, we show you how to edit the snapshot schedule of a volume.

### To edit the snapshot schedule for a volume

1. On the AWS Storage Gateway console, select the gateway that contains the volume snapshot schedule that you want to edit.
2. Click the **Snapshot Schedules** tab.

   The tab shows a list of your storage volumes on the selected gateway.

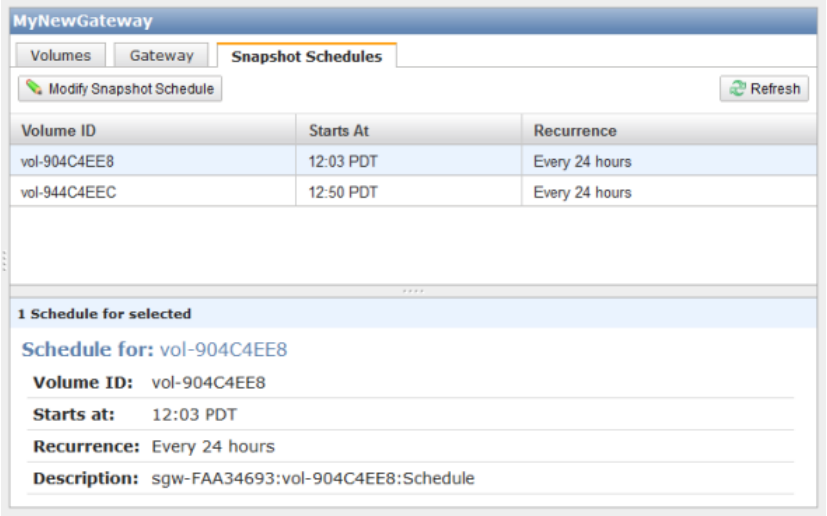


3. Select a volume.

   The AWS Storage Gateway console shows the snapshot schedule details for this volume.
4. Click **Modify Snapshot Schedule**.

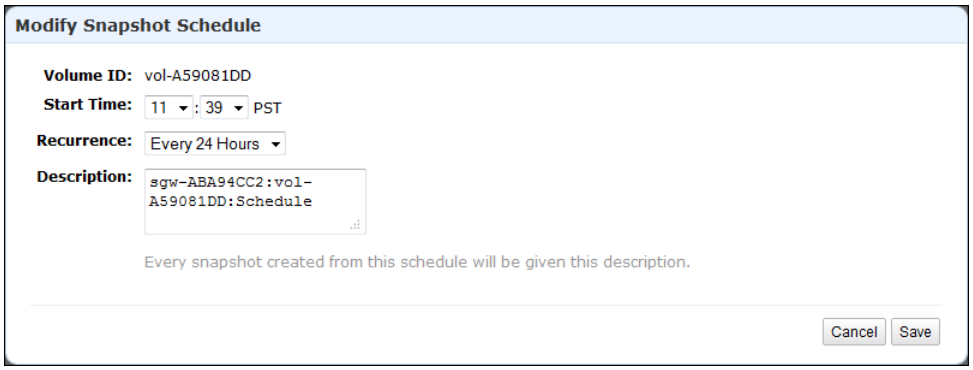


5. In the **Modify Snapshot Schedule** dialog box, update the schedule fields as needed. For example, you can increase the default snapshot frequency of once a day or change the time a snapshot is taken.
6. Click **Save** to save the snapshot schedule updates.

# Creating an Ad Hoc Snapshot

In addition to scheduled snapshots, AWS Storage Gateway allows you to take one-time, ad hoc snapshots. Doing this lets you back up your storage volume immediately without waiting for the next scheduled snapshot.

### To take an ad hoc snapshot of your storage volume

1. On the AWS Storage Gateway console, select the gateway that contains the storage volume of which you want to take a snapshot.
2. Click the **Volumes** tab.
3. Select a volume from the list, and then click **Create Snapshot**.

   AWS Storage Gateway starts the snapshot process immediately.

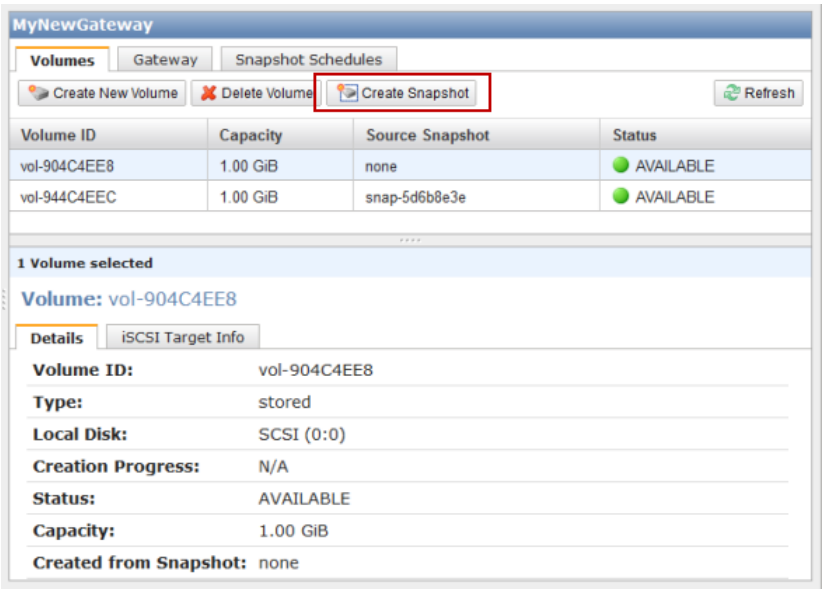


4. Verify the snapshot using the console. For more information, see Finding a Snapshot (p. 203).

# Deleting a Snapshot

You might want to delete a snapshot, for example, if you have taken many snapshots of a storage volume over a period of time and you don't need the older snapshots. Because snapshots are incremental backups, if you delete a snapshot, only the data that is not needed in other snapshots is deleted.

### Topics

- Deleting Snapshots Using the AWS SDK for Java (p. 213)
- Deleting Snapshots Using the AWS SDK for .NET (p. 217)
- Deleting Snapshots Using the AWS Tools for Windows PowerShell (p. 220)

On the AWS Storage Gateway console, you can delete a snapshots one at a time. To delete many snapshots, you can use one of the AWS SDKs that supports AWS Storage Gateway operations. For examples, see Deleting Snapshots Using the AWS SDK for Java (p. 213), Deleting Snapshots Using the AWS SDK for .NET (p. 217), and Deleting Snapshots Using the AWS Tools for Windows PowerShell (p. 220).

### To delete a snapshot using the console

1. On the AWS Storage Gateway console, click **Snapshots** in the **Navigation** pane.

   A list of snapshots appears in the main pane.

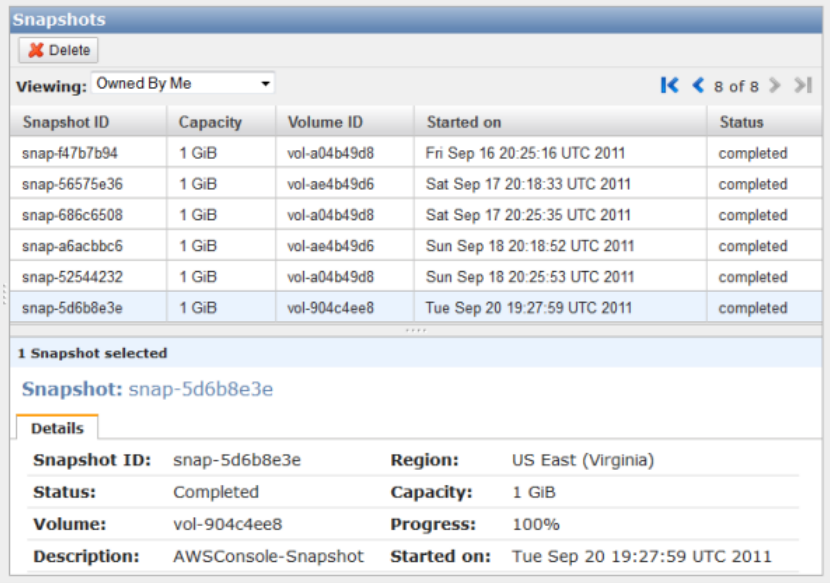
Snapshots
Delete
Viewing: Owned By Me
8 of 8
Snapshot ID | Capacity | Volume ID | Started on | Status
snap-f47b7b94 | 1 GiB | vol-a04b49d8 | Fri Sep 16 20:25:16 UTC 2011 | completed
snap-56575e36 | 1 GiB | vol-ae4b49d6 | Sat Sep 17 20:18:33 UTC 2011 | completed
snap-686c6508 | 1 GiB | vol-a04b49d8 | Sat Sep 17 20:25:35 UTC 2011 | completed
snap-a6acbbc6 | 1 GiB | vol-ae4b49d6 | Sun Sep 18 20:18:52 UTC 2011 | completed
snap-52544232 | 1 GiB | vol-a04b49d8 | Sun Sep 18 20:25:53 UTC 2011 | completed
snap-5d6b8e3e | 1 GiB | vol-904c4ee8 | Tue Sep 20 19:27:59 UTC 2011 | completed
1 Snapshot selected
Snapshot: snap-5d6b8e3e
Details
Snapshot ID: snap-5d6b8e3e
Region: US East (Virginia)
Status: Completed
Capacity: 1 GiB
Volume: vol-904c4ee8
Progress: 100%
Description: AWSConsole-Snapshot
Started on: Tue Sep 20 19:27:59 UTC 2011

2. Select the snapshot that you want to delete and click **Delete**.

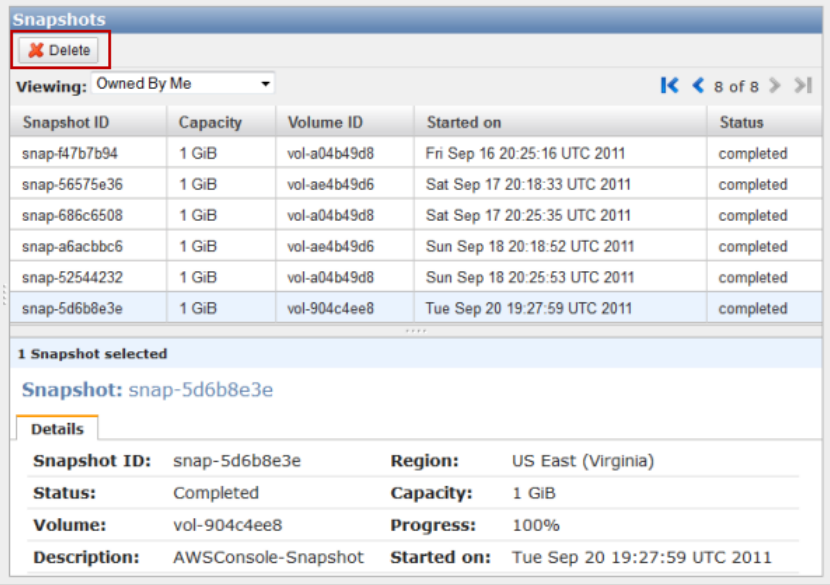
Snapshots
Delete
Viewing: Owned By Me
8 of 8
Snapshot ID | Capacity | Volume ID | Started on | Status
snap-f47b7b94 | 1 GiB | vol-a04b49d8 | Fri Sep 16 20:25:16 UTC 2011 | completed
snap-56575e36 | 1 GiB | vol-ae4b49d6 | Sat Sep 17 20:18:33 UTC 2011 | completed
snap-686c6508 | 1 GiB | vol-a04b49d8 | Sat Sep 17 20:25:35 UTC 2011 | completed
snap-a6acbbc6 | 1 GiB | vol-ae4b49d6 | Sun Sep 18 20:18:52 UTC 2011 | completed
snap-52544232 | 1 GiB | vol-a04b49d8 | Sun Sep 18 20:25:53 UTC 2011 | completed
snap-5d6b8e3e | 1 GiB | vol-904c4ee8 | Tue Sep 20 19:27:59 UTC 2011 | completed
1 Snapshot selected
Snapshot: snap-5d6b8e3e
Details
Snapshot ID: snap-5d6b8e3e
Region: US East (Virginia)
Status: Completed
Capacity: 1 GiB
Volume: vol-904c4ee8
Progress: 100%
Description: AWSConsole-Snapshot
Started on: Tue Sep 20 19:27:59 UTC 2011

3. To confirm that you want to delete the snapshot, click **OK**.

## Deleting Snapshots Using the AWS SDK for Java

To delete many snapshots associated with a volume, you can use a programmatic approach. The example following demonstrates how to delete snapshots using the AWS SDK for Java. To use the example code, you should be familiar with running a Java console application. For more information, go to Getting Started

in the *AWS SDK for Java Developer Guide*. If you need to just delete a few snapshots, use the console as described in To delete a snapshot using the console (p. 213).

### Example : Deleting Snapshots Using the AWS SDK for Java

The following Java code example lists the snapshots for each volume of a gateway and whether the snapshot start time is before or after a specified date. It uses the AWS SDK for Java API for AWS Storage Gateway and Amazon EC2. The Amazon EC2 API includes operations for working with snapshots.

You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the number of days back you want to save snapshots (snapshots taken before this cutoff are deleted). You also need to specify the Boolean value `viewOnly`, which indicates whether you want to view the snapshots that would be deleted or actually perform the snapshot deletions. You should run the code first with just the view option (that is, with `viewOnly` set to `true`) to see what the code would delete. For a list of AWS service endpoints you can use with AWS Storage Gateway, go to Regions and Endpoints in the *Amazon Web Services Glossary*.

```
import java.io.IOException;
import java.util.ArrayList;
import java.util.Calendar;
import java.util.Collection;
import java.util.Date;
import java.util.GregorianCalendar;
import java.util.List;

import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.ec2.AmazonEC2Client;
import com.amazonaws.services.ec2.model.DeleteSnapshotRequest;
import com.amazonaws.services.ec2.model.DescribeSnapshotsRequest;
import com.amazonaws.services.ec2.model.DescribeSnapshotsResult;
import com.amazonaws.services.ec2.model.Filter;
import com.amazonaws.services.ec2.model.Snapshot;
import com.amazonaws.services.storagegateway.AWSStorageGatewayClient;
import com.amazonaws.services.storagegateway.model.ListVolumesRequest;
import com.amazonaws.services.storagegateway.model.ListVolumesResult;
import com.amazonaws.services.storagegateway.model.VolumeInfo;

public class ListDeleteVolumeSnapshotsExample {

    public static AWSStorageGatewayClient sgClient;
    public static AmazonEC2Client ec2Client;
    static String serviceURLSG = "https://storagegateway.us-east-1.amazon
aws.com";
    static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";

    // The gatewayARN
    public static String gatewayARN = "*** provide gateway ARN ***";

    // The number of days back you want to save snapshots. Snapshots before
this cutoff are deleted
    // if viewOnly = false.
    public static int daysBack = 10;

   // true = show what will be deleted; false = actually delete snapshots that
 meet the daysBack criteria
    public static boolean viewOnly = true;

    public static void main(String[] args) throws IOException {

        // Create a storage gateway and amazon ec2 client
        sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(
```

```
                ListDeleteVolumeSnapshotsExample.class.getResourceAs
Stream("AwsCredentials.properties")));
        sgClient.setEndpoint(serviceURLSG);

        ec2Client = new AmazonEC2Client(new PropertiesCredentials(
                ListDeleteVolumeSnapshotsExample.class.getResourceAs
Stream("AwsCredentials.properties")));
        ec2Client.setEndpoint(serviceURLEC2);

        List<VolumeInfo> volumes = ListVolumesForGateway();
        DeleteSnapshotsForVolumes(volumes, daysBack);

    }
    public static List<VolumeInfo> ListVolumesForGateway()
    {
        List<VolumeInfo> volumes = new ArrayList<VolumeInfo>();

        String marker = null;
        do {
            ListVolumesRequest request = new ListVolumesRequest().withGatewa
yARN(gatewayARN);
            ListVolumesResult result = sgClient.listVolumes(request);
            marker = result.getMarker();

            for (VolumeInfo vi : result.getVolumeInfos())
            {
                volumes.add(vi);
                System.out.println(OutputVolumeInfo(vi));
            }
        } while (marker != null);

        return volumes;
    }
    private static void DeleteSnapshotsForVolumes(List<VolumeInfo> volumes,
            int daysBack2) {

        // Find snapshots and delete for each volume
        for (VolumeInfo vi : volumes) {

            String volumeARN = vi.getVolumeARN();
            String volumeId = volumeARN.substring(volumeARN.lastIndex
Of("/")+1).toLowerCase();
            Collection<Filter> filters = new ArrayList<Filter>();
            Filter filter = new Filter().withName("volume-id").withVal
ues(volumeId);
            filters.add(filter);

            DescribeSnapshotsRequest describeSnapshotsRequest =
                new DescribeSnapshotsRequest().withFilters(filters);
            DescribeSnapshotsResult describeSnapshotsResult =
                ec2Client.describeSnapshots(describeSnapshotsRequest);

            List<Snapshot> snapshots = describeSnapshotsResult.getSnapshots();

            System.out.println("volume-id = " + volumeId);
            for (Snapshot s : snapshots){
                StringBuilder sb = new StringBuilder();
                boolean meetsCriteria = !CompareDates(daysBack, s.getStart
```

```
Time());
                sb.append(s.getSnapshotId() + ", " + s.getStartTime().to
String());
                sb.append(", meets criteria for delete? " + meetsCriteria);
                sb.append(", deleted? ");
                if (!viewOnly & meetsCriteria) {
                    sb.append("yes");
                    DeleteSnapshotRequest deleteSnapshotRequest =
                        new DeleteSnapshotRequest().withSnapshotId(s.getSnap
shotId());
                    ec2Client.deleteSnapshot(deleteSnapshotRequest);
                }
                else {
                    sb.append("no");
                }
                System.out.println(sb.toString());
            }
        }
    }

    private static String OutputVolumeInfo(VolumeInfo vi) {

        String volumeInfo = String.format(
                 "Volume Info:\n" +
                 "  ARN: %s\n" +
                 "  Type: %s\n",
                 vi.getVolumeARN(),
                 vi.getVolumeType());
        return volumeInfo;
     }

    // Returns the date in two formats as a list
    public static boolean CompareDates(int daysBack, Date snapshotDate) {
        Date today = new Date();
        Calendar cal = new GregorianCalendar();
        cal.setTime(today);
        cal.add(Calendar.DAY_OF_MONTH, -daysBack);
        Date cutoffDate = cal.getTime();
        return (snapshotDate.compareTo(cutoffDate) > 0) ? true : false;
    }

}
```

## Deleting Snapshots Using the AWS SDK for .NET

To delete many snapshots associated with a volume, you can use a programmatic approach. The example following demonstrates how to delete snapshots using the AWS SDK for .NET. To use the example code, you should be familiar with running a .NET console application. For more information, go to Getting Started in the *AWS SDK for .NET Developer Guide*. If you need to just delete a few snapshots, use the console as described in To delete a snapshot using the console (p. 213).

### Example : Deleting Snapshots Using the AWS SDK for .NET

The following C# code example lists the snapshots for each volume of a gateway and whether the snapshot start time is before or after a specified date. It uses the AWS SDK for .NET API for AWS Storage Gateway and Amazon EC2. The Amazon EC2 API includes operations for working with snapshots.

> **Note**
> The following code example uses the AWS SDK for .NET version 1 and might include methods that are no longer available in the current version of the SDK. For more information, see Migrating Your Code to the Latest Version of the AWS SDK for .NET.

You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the number of days back you want to save snapshots (snapshots taken before this cutoff are deleted). You also need to specify the Boolean value `viewOnly`, which indicates whether you want to view the snapshots that would be deleted or actually perform the snapshot deletions. You should run the code first with just the view option (that is, with `viewOnly` set to `true`) to see what the code would delete. For a list of AWS service endpoints you can use with AWS Storage Gateway, go to Regions and Endpoints in the *Amazon Web Services Glossary*.

```
using System;
using System.Collections.Generic;
using System.Text;
using Amazon.StorageGateway.Model;
using Amazon.StorageGateway;
using Amazon.EC2;
using Amazon.EC2.Model;

namespace AWSStorageGateway
{
    class ListDeleteVolumeSnapshotsExample
    {
        static AmazonStorageGatewayClient sgClient;
        static AmazonStorageGatewayConfig sgConfig;
        static AmazonEC2Config ec2Config;
        static AmazonEC2Client ec2Client;
        static String serviceURLSG = "https://storagegateway.us-east-1.amazon
aws.com";
        static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";

        // The gatewayARN
        public static String gatewayARN = "*** provide gateway ARN ***";

       // The number of days back you want to save snapshots. Snapshots before
 this cutoff are deleted
        // if viewOnly = false.
        public static int daysBack = 10;

        // true = show what will be deleted; false = actually delete snapshots
 that meet the daysBack criteria
        public static Boolean viewOnly = true;

        public static void Main(string[] args)
        {
            // Create a storage gateway client
            sgConfig = new AmazonStorageGatewayConfig();
            sgConfig.ServiceURL = serviceURLSG;
            sgClient = new AmazonStorageGatewayClient(sgConfig);
```

```
            //Create a ec2 client
            ec2Config = new AmazonEC2Config();
            ec2Config.ServiceURL = serviceURLEC2;
            ec2Client = new AmazonEC2Client(ec2Config);

            List<VolumeInfo> volumes = ListVolumesForGateway();
            DeleteSnapshotsForVolumes(volumes, daysBack);

            Console.WriteLine("\nTo continue, press Enter.");
            Console.Read();
        }
        public static List<VolumeInfo> ListVolumesForGateway()
        {
            ListVolumesRequest request = new ListVolumesRequest().WithGatewa
yARN(gatewayARN);
            ListVolumesResponse response = sgClient.ListVolumes(request);
            ListVolumesResult volumes = response.ListVolumesResult;

            foreach (VolumeInfo vi in volumes.VolumeInfos)
            {
                Console.WriteLine(OutputVolumeInfo(vi));
            }

            return volumes.VolumeInfos;
        }

        public static void DeleteSnapshotsForVolumes(List<VolumeInfo> volumes,
 int cutoffDays)
        {

            foreach (VolumeInfo vi in volumes)
            {
                String volumeARN = vi.VolumeARN;
              String volumeId = volumeARN.Substring(volumeARN.LastIndexOf("/")
 + 1).ToLower();
                Filter[] filters = new Filter[1];
                filters[0] = new Filter().WithName("volume-id").With
Value(volumeId);

                DescribeSnapshotsRequest describeSnapshotsRequest =
                    new DescribeSnapshotsRequest().WithFilter(filters);
                DescribeSnapshotsResponse describeSnapshotsResponse =
                    ec2Client.DescribeSnapshots(describeSnapshotsRequest);

                List<Snapshot> snapshots = describeSnapshotsResponse.Describe
SnapshotsResult.Snapshot;
                Console.WriteLine("volume-id = " + volumeId);
                foreach (Snapshot s in snapshots)
                {
                    StringBuilder sb = new StringBuilder();
                  Boolean meetsCriteria = CompareDates(daysBack, s.StartTime);

                    sb.Append(s.SnapshotId + ", " + s.StartTime);
                   sb.Append(", meets criteria for delete? " + meetsCriteria);

                    sb.Append(", deleted? ");
                    if (!viewOnly & meetsCriteria)
                    {
```

```
                        sb.Append("yes");
                        DeleteSnapshotRequest deleteSnapshotRequest =
                            new DeleteSnapshotRequest().WithSnapshotId(s.Snap
shotId);
                        ec2Client.DeleteSnapshot(deleteSnapshotRequest);
                    }
                    else
                    {
                        sb.Append("no");
                    }
                    Console.WriteLine(sb.ToString());
                }
            }

        }

        private static String OutputVolumeInfo(VolumeInfo vi)
        {

            String volumeInfo = String.Format(
                     "Volume Info:\n" +
                     "  ARN: {0}\n" +
                     "  Type: {1}\n",
                     vi.VolumeARN,
                     vi.VolumeType);
            return volumeInfo;
        }

        public static Boolean CompareDates(int daysBack, String d)
        {
            DateTime snapshotDate = DateTime.Parse(d);
            DateTime cutoffDate = DateTime.Now.Add(new TimeSpan(-daysBack, 0,
0, 0));
            return (DateTime.Compare(snapshotDate, cutoffDate) < 0) ? true :
false;
        }

    }
}
```

## Deleting Snapshots Using the AWS Tools for Windows PowerShell

To delete many snapshots associated with a volume, you can use a programmatic approach. The example following demonstrates how to delete snapshots using the AWS Tools for Windows PowerShell. To use the example script, you should be familiar with running a PowerShell script. For more information, go to Getting Started in the *AWS Tools for Windows PowerShell*. If you need to delete just a few snapshots, use the console as described in To delete a snapshot using the console (p. 213).

### Example : Deleting Snapshots Using the AWS Tools for Windows PowerShell

The following PowerShell script example lists the snapshots for each volume of a gateway and whether the snapshot start time is before or after a specified date. It uses the AWS Tools for Windows PowerShell cmdlets for AWS Storage Gateway and Amazon EC2. The Amazon EC2 API includes operations for working with snapshots.

You need to update the script and provide your gateway Amazon Resource Name (ARN) and the number of days back you want to save snapshots (snapshots taken before this cutoff are deleted). You also need to specify the Boolean value `viewOnly`, which indicates whether you want to view the snapshots that would be deleted or actually perform the snapshot deletions. You should run the code first with just the view option (that is, with `viewOnly` set to `true`) to see what the code would delete.

```
<#
.DESCRIPTION
    Delete Snapshots of a specified volume that match given criteria.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from http://console.amazonaws.cn/powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
   For more info see, http://docs.amazonaws.cn/powershell/latest/userguide//spe
cifying-your-aws-credentials.html

.EXAMPLE
    powershell.exe .\SG_DeleteSnapshots.ps1
#>

# Criteria to use to filter the results returned.
$daysBack = 18
$gatewayARN = "*** provide gateway ARN ***"
$viewOnly = $true;

#ListVolumes
$volumesResult = Get-SGVolume -GatewayARN $gatewayARN
$volumes = $volumesResult.VolumeInfos
Write-Output("`nVolume List")
foreach ($volumes in $volumesResult)
  { Write-Output("`nVolume Info:")
    Write-Output("ARN:  " + $volumes.VolumeARN)
    write-Output("Type: " + $volumes.VolumeType)
  }

Write-Output("`nWhich snapshots meet the criteria?")
foreach ($volume in $volumesResult)
  {
    $volumeARN = $volume.VolumeARN
    $volumeId = $volumeARN.Substring($volumeARN.LastIndexOf("/")+1).ToLower()

    $filter = New-Object Amazon.EC2.Model.Filter
    $filter.Name = "volume-id"
    $filter.Value.Add($volumeId)

    $snapshots = get-EC2Snapshot -Filter $filter
    Write-Output("`nFor volume-id = " + $volumeId)
    foreach ($s in $snapshots)
    {
       $d = ([DateTime]::Now).AddDays(-$daysBack)
```

```
      $meetsCriteria = $false
      if ([DateTime]::Compare($d, $s.StartTime) -gt 0)
      {
           $meetsCriteria = $true
      }

     $sb = $s.SnapshotId + ", " + $s.StartTime + ", meets criteria for delete?
" + $meetsCriteria
      if (!$ViewOnly -AND $meetsCriteria)
      {
          $resp = Remove-EC2Snapshot -SnapshotId $s.SnapshotId
          #Can get RequestId from response for troubleshooting.
          $sb = $sb + ", deleted? yes"
      }
      else {
          $sb = $sb + ", deleted? no"
      }
      Write-Output($sb)
   }
 }
```

# Restoring a Snapshot

You can restore a snapshot of a volume to a new AWS Storage Gateway volume, or you can use the snapshot to create an Amazon EBS volume and attach this volume to an Amazon EC2 instance. When you restore the snapshot to a new AWS Storage Gateway volume, you can mount the volume as an iSCSI device to your on-premises application server. You can then access the contents of the snapshot. Overall, the process is similar to creating a new volume.

The use cases for restoring snapshots depends on the type of gateway you set up—for more information on gateway types, see How AWS Storage Gateway Works (Architecture) (p. 2).

- For gateway-cached volumes, your volume data is already stored in Amazon S3. In this case, snapshots are typically used to preserve older versions of your data. After initiating a snapshot restore to a gateway-cached volume, snapshot data is downloaded to the local cache only upon first access of the data.
- For gateway-stored volumes, your volume data is stored on-premises, In this case, snapshots provide durable, off-site backups in Amazon S3. For example, if a local disk allocated as a storage volume crashes, you can provision a new local disk and restore a snapshot to it during the volume creation process. For more information on this approach, see Adding a Storage Volume (p. 100).

  After you initiate a snapshot restore to a gateway-stored volume, snapshot data is downloaded in the background. This functionality means that once you create a volume from a snapshot, there is no need to wait for all of the data to transfer from Amazon S3 to your volume before your application can start accessing the volume and all of its data. If your application accesses a piece of data that has yet to be loaded, the gateway immediately downloads the requested data from Amazon S3. The gateway then continues loading the rest of the volume's data in the background.

**Topics**


- Restoring a Snapshot to a Storage Volume (p. 222)
- Restoring a Snapshot to an Amazon EBS Volume (p. 225)


## Restoring a Snapshot to a Storage Volume

The following procedure applies to both gateway-cached and gateway-stored volumes.

### To create a storage volume from an existing snapshot

1. On the AWS Storage Gateway console, click **Snapshots**.

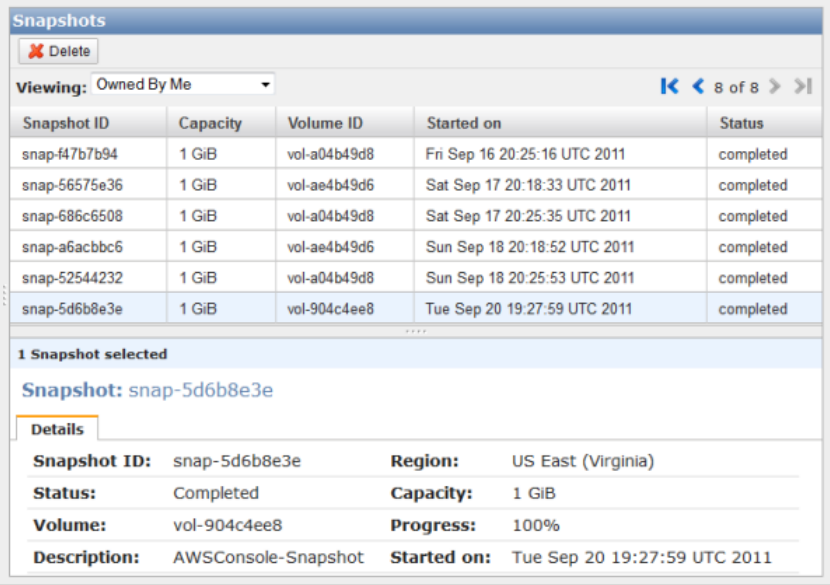


2. In the snapshot list, select the snapshot you want to create a storage volume from, and then note its **Snapshot ID** value for use in a later step.

   **Important**

   If the Amazon EBS volume you create the snapshot from contains data from a public Amazon Machine Image (AMI), such as an AMI from AWS Marketplace, you will not have the permission to create the volume.
   Storage volumes cannot be created from snapshots that are encrypted.
   The size of the storage volume you create must be equal to or greater than the size of the snapshot you create it from. For information on how to add a disk to your gateway VM that can be used as a gateway-stored volume, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23). Once you add such a volume, you can access the contents of the volume from your on-premises applications. For more information, see Connecting to Storage Volumes (p. 111).

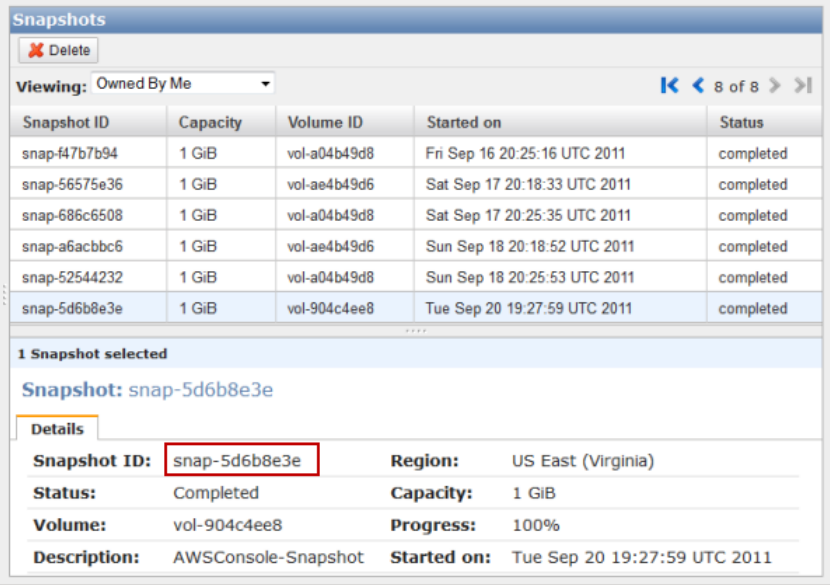


3. In the **Navigation** pane, select the gateway to which you want to restore the snapshot.
4. Click **Create Volume**.

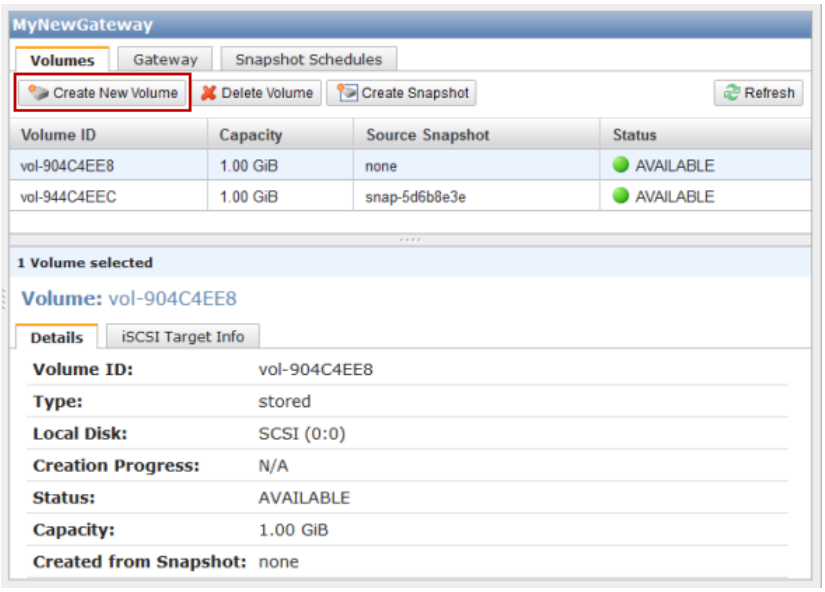


5. Depending on the type of gateway you configured, choose one of the following steps.

   a. To create a new gateway-stored volume, do the following steps.

      i. In the **Based on Snapshot ID** field of the **Create Storage Volume** dialog box, paste the snapshot ID that you copied previously.

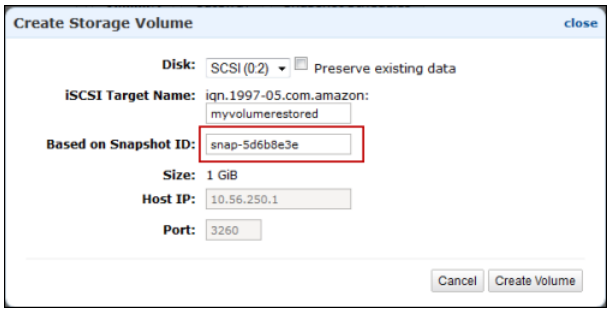


ii. Select a disk, type in a unique volume target name, and then click **Create Volume**.

> **Note**
> If you select **Preserve existing data**, the volume creation process could take a long time to complete. The maximuum total capacity of a volume is 192 TBs.

b. To create a new gateway-cached volume, do the following steps.

i. In the **Based on Snapshot ID** field of the **Configure Your Gateway** dialog box, paste the snapshot ID you copied previously.

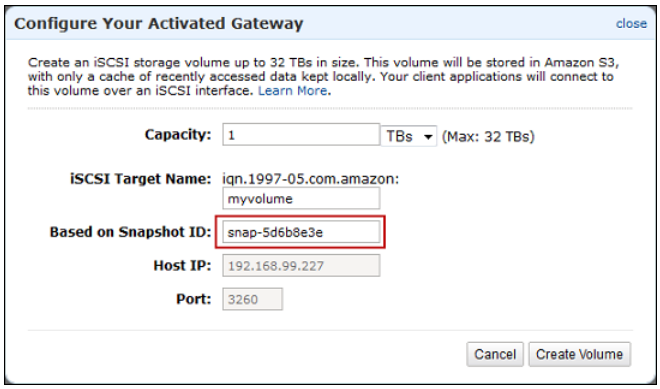
Configure Your Activated Gateway
close
Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3, with only a cache of recently accessed data kept locally. Your client applications will connect to this volume over an iSCSI interface. Learn More.
Capacity:
1
TBs
(Max: 32 TBs)
iSCSI Target Name:
iqn.1997-05.com.amazon:
myvolume
Based on Snapshot ID:
snap-5d6b8e3e
Host IP:
192.168.99.227
Port:
3260
Cancel
Create Volume

ii. Select a capacity for the disk, type in a unique volume target name, and then click **Create Volume**.

The size of your storage volume must be greater than or equal to the size of the snapshot. You can now access the contents of this volume from your on-premises applications. For more information on this process, see Connecting to Storage Volumes (p. 111).

## Restoring a Snapshot to an Amazon EBS Volume

Your snapshots of local storage volumes taken by AWS Storage Gateway are stored in Amazon S3 as Amazon EBS snapshots. For snapshots up to 16 TiB in size, you can restore snapshots of your local storage volumes to an Amazon EBS volume. You can then attach the Amazon EBS volume to an Amazon EC2 instance. Doing this allows you to easily migrate data from your on-premises applications to your applications running on Amazon EC2, in case you need to use the Amazon EC2 compute capacity for disaster recovery or data processing. To see detailed pricing for Amazon EC2 and Amazon EBS, go to the Amazon EC2 Pricing page.

**To restore a snapshot to an Amazon EBS volume**

1. Create an Amazon EBS volume.

   - Follow the instructions in Creating an Amazon EBS Volume in the *Amazon Elastic Compute Cloud User Guide*.

     The volume size that you specify must be greater than or equal to the size of the snapshot. To specify the snapshot to use, select its snapshot ID in the drop-down list of the **Create Volume** wizard in the **EBS Volumes** pane of the Amazon EC2 console. Alternatively, you can use the Amazon EC2 API to create your Amazon EBS volumes.

     **Note**
     By default, AWS limits the maximum size of Amazon EBS volumes you can create per AWS account. For information about default Amazon EBS volume limits, see Amazon EBS Limits. in the *Amazon EC2 User Guide*. For information about how to increase these limits, go to Amazon EC2 Service Limits.

2. Attach the Amazon EBS volume to an Amazon EC2 instance. For more information, go to Attaching the Volume to an Instance in the *Amazon Elastic Compute Cloud User Guide*.

# Troubleshooting Your Volume Gateway (Gateway-Cached and Gateway-Stored)

This section discusses troubleshooting issues related to gateways and storage volumes. The gateway troubleshooting issues are split into two sections: gateways that are on-premises and gateways that are deployed on Amazon EC2. The on-premises gateway troubleshooting issues cover gateways deployed on both the VMware ESXi and Microsoft Hyper-V clients. The troubleshooting issues for storage volumes apply to all gateway types.

**Topics**

- Troubleshooting On-Premises Gateway Issues (p. 226)
- Troubleshooting Amazon EC2 Gateway Issues (p. 230)
- Troubleshooting Storage Volume Issues (p. 233)
- Troubleshooting Snapshot Issues (p. 237)
- Using Recovery Snapshots (Gateway-Cached) (p. 237)

## Troubleshooting On-Premises Gateway Issues

The following table lists typical issues that you might encounter working with your on-premises gateways.

| Issue | Action to Take |
| --- | --- |
| You cannot find the IP address of your gateway. | Use the hypervisor client to connect to your host to find the gateway IP address.<br>• For VMware ESXi, the VM's IP address can be found in the vSphere client on the **Summary** tab. For more information, see Step 2.4: Activate the AWS Storage Gateway (p. 30).<br>• For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console. For more information, see Step 2.4: Activate Your Storage Gateway (p. 50).<br><br>If you are still having trouble finding the gateway IP address:<br>• Check that the VM is turned on. Only when the VM is turned on does an IP address get assigned to your gateway.<br>• Wait for the VM to finish startup. If you just turned on your VM, then it might take several minutes for the gateway to finish its boot sequence. |
| Your gateway's activation fails when you click the **Proceed to Activation** button in the AWS Storage Gateway console. | • Check that the gateway VM can be accessed by pinging the VM from your client.<br>• Check that your VM has network connectivity to the Internet. Otherwise, you'll need to configure a SOCKS proxy. For more information on doing so, see Routing Your On-Premises Gateway Through a Proxy (p. 141).<br>• Check that the host has the correct time, that the host is configured to synchronize its time automatically to a Network Time Protocol (NTP) server, and that the gateway VM has the correct time. For information about synchronizing the time of hypervisor hosts and VMs, see Synchronizing Your Gateway VM Time (p. 146).<br>• After performing these steps, you can retry the gateway deployment using the AWS Storage Gateway console and the **Setup and Activate Gateway** wizard.<br>• Check that your VM has at least 7.5 GB of RAM. Gateway allocation fails if there is less than 7.5 GB of RAM. For more information, see Requirements (p. 11). |
| You need to remove a disk allocated as upload buffer space. For example, you might want to reduce the amount of upload buffer space for a gateway, or you might need to replace a disk used as an upload buffer that has failed. | For instructions about removing a disk allocated as upload buffer space, see Removing Upload Buffer Capacity (p. 173) or Removing Upload Buffer Capacity (Gateway-Stored) (p. 181). |

| Issue | Action to Take |
| --- | --- |
| You need to improve bandwidth between your gateway and AWS. | You can improve the bandwidth from your gateway to AWS by setting up your Internet connection to AWS on a network adapter (NIC) separate from that connecting your applications and the gateway VM. Taking this approach is useful if you have a high-bandwidth connection to AWS and you want to avoid bandwidth contention, especially during a snapshot restore. For high-throughput workload needs, you can use AWS Direct Connect to establish a dedicated network connection between your on-premises gateway and AWS. To measure the bandwidth of the connection from your gateway to AWS, use the `CloudBytesDownloaded` and `CloudBytesUploaded` metrics of the gateway. For more on this subject, see Measuring Performance Between Your Gateway and AWS (p. 188). Improving your Internet connectivity helps to ensure that your upload buffer does not fill up. |
| Throughput to or from your gateway drops to zero. | • On the **Gateway** tab of the AWS Storage Gateway console, verify that the IP addresses for your gateway VM are the same that you see using your hypervisor client software (that is, the VMware Vsphere client or Microsoft Hyper-V Manager). If you find a mismatch, restart your gateway from the AWS Storage Gateway console, as shown in Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console (p. 127). After the restart, the addresses in the **IP Addresses** list in the AWS Storage Gateway console's **Gateway** tab should match the IP addresses for your gateway, which you determine from the hypervisor client.<br>  • For VMware ESXi, the VM's IP address can be found in the vSphere client on the **Summary** tab. For more information, see Step 2.4: Activate the AWS Storage Gateway (p. 30).<br>  • For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console. For more information, see Step 2.4: Activate Your Storage Gateway (p. 50).<br>• Check your gateway's connectivity to AWS as described in Testing Your Gateway Connection to the Internet (p. 145).<br>• Check your gateway's network adapter configuration, and ensure that all the interfaces you intended to be enabled for the gateway are enabled. To view the network adapter configuration for your gateway, follow the instructions in Configuring Your Gateway to Use a Static IP Address (p. 142) and select the option for viewing your gateway's network configuration.<br><br>You can view the throughput to and from your gateway from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway to AWS, see Measuring Performance Between Your Gateway and AWS (p. 188). |
| You are having trouble importing (deploying) AWS Storage Gateway on Microsoft Hyper-V. | See Troubleshooting Your Microsoft Hyper-V Setup (p. 417), which discusses some of the common issues of deploying a gateway on Microsoft Hyper-V. |

## Allowing AWS Storage Gateway Access through Firewalls and Routers

Your locally deployed gateway requires access to the following endpoints to communicate with AWS. If you use firewalls or routers to filter or limit network traffic, you must configure your firewall and router to allow these service endpoints for outbound communication to AWS.

```
client-cp.storagegateway.region.amazonaws.com:443
dp-1.storagegateway.region.amazonaws.com:443
anon-cp.storagegateway.region.amazonaws.com:443
proxy-app.storagegateway.region.amazonaws.com:443
storagegateway.region.amazonaws.com:443
```

The following table provides a list of region strings for the available regions.

| Region Name | Region String |
|---|---|
| US East (N. Virginia) | `us-east-1` |
| US West (N. California) | `us-west-1` |
| US West (Oregon) | `us-west-2` |
| EU (Ireland) | `eu-west-1` |
| EU (Frankfurt) | `eu-central-1` |
| Asia Pacific (Singapore) | `ap-southeast-1` |
| Asia Pacific (Sydney) | `ap-southeast-2` |
| Asia Pacific (Tokyo) | `ap-northeast-1` |
| South America (Sao Paulo) | `sa-east-1` |

Depending on your gateway's region, you replace *`region`* in the endpoint with the corresponding region string. For example, if you create a gateway in the US West (Oregon) region, the endpoint looks like this: `storagegateway.us-west-2.amazonaws.com:443`.

## Enabling AWS Support Access to Your Gateway Hosted On-Premises

AWS Storage Gateway provides a local console you can use to perform several maintenance tasks, including enabling AWS Support to access your gateway to assist you with troubleshooting gateway issues. By default, AWS Support's access to your gateway is disabled. You enable this access through the host's local console. To give AWS Support access to your gateway, you first log in to the local console for the host, navigate to the storage gateway's console, and then connect to the support server.

### To enable AWS Support access to your gateway

1. Log in to your host's local console. For instructions, go to Logging Into Your Gateway Local Console.

   The local console looks like the following.

2. At the prompt, type **5** to open the AWS Storage Gateway console.

3. Type **h** to open the **AVAILABLE COMMANDS** window.
4. In the **AVAILABLE COMMANDS** window, type **`open-support-channel`** to connect to customer support for AWS Storage Gateway.
5. At the prompt, type a support service number for your gateway to connect to the AWS Storage Gateway support server. Choose any number between 2000 and 65000.

Once the support channel is established, provide the support service number you chose to AWS Support so AWS Support can provide troubleshooting assistance.

6. When the support session is completed, type **q** to end it.
7. Type **`exit`** to log out of the AWS Storage Gateway console.
8. Follow the prompts to exit the local console.

# Troubleshooting Amazon EC2 Gateway Issues

The following table lists typical issues that you might encounter working with your gateway deployed on Amazon Elastic Compute Cloud (Amazon EC2). For more information about the difference between an on-premises gateway and a gateway deployed in Amazon EC2, see Amazon EC2 Gateway (p. 57).

| Issue | Action to Take |
| --- | --- |
| Your Amazon EC2 gateway activation fails when you click the **Proceed to Activation** button in the AWS Storage Gateway console. | If activation has not occurred in a few moments, check the following in the Amazon EC2 console:<br>• Port 80 is enabled in the security group you associated with the instance. For more information about adding a security group rule, go to Adding a Security Group Rule in the *Amazon EC2 User Guide for Linux Instances*.<br>• The gateway instance is marked as running. In the Amazon EC2 console, the **State** of the instance should be RUNNING.<br><br>After correcting the problem, try activating the gateway again by going to the AWS Storage Gateway console, clicking **Deploy a new Gateway on Amazon EC2**, and re-entering the IP address of the instance. |
| You can't find your Amazon EC2 gateway instance in the list of instances. | If you did not give your instance a resource tag and you have many instances running, it can be hard to tell which instance you deployed the gateway in. In this case, you can take the following actions to find the gateway instance:<br>• Check the name of the Amazon Machine Image (AMI) name on the **Description** tab of the instance. An instance based on the AWS Storage Gateway AMI should start with the text **`aws-storage-gateway-ami`**.<br>• If you have several instances based on the AWS Storage Gateway AMI, check the instance launch time to find the correct instance. |
| You created an Amazon EBS volume but can't attach it to your Amazon EC2 gateway instance. | Check that the Amazon EBS volume in question is in the same Availability Zone as the gateway instance. If there is a discrepancy in Availability Zones, create a new Amazon EBS volume in the same Availability Zone as your instance. |
| You can't attach an initiator to a storage volume target of your Amazon EC2 gateway. | Check that the security group you launched the instance with includes a rule allowing the port that you are using for iSCSI access. The port is usually set as 3260. For more information on connecting to storage volumes, see Connecting to Storage Volumes (p. 111). |
| You activated your Amazon EC2 gateway, but when you try to add storage volumes, you receive an error message indicating you have no disks available. | For a newly activated gateway, no volume storage is defined. Before you can define volume storage, you must allocate local disks to the gateway to use as an upload buffer and cache storage. For a gateway deployed to Amazon EC2, the local disks are Amazon EBS volumes attached to the instance. This error message likely occurs because no Amazon EBS volumes are defined for the instance.<br><br>Check block devices defined for the instance that is running the gateway. If there are only two block devices (the default devices that come with the AMI), then you should add storage. For more information on doing so, see Amazon EC2 Gateway (p. 57). After attaching two or more Amazon EBS volumes, try creating volume storage on the gateway. |
| You need to remove a disk allocated as upload buffer space because you want to reduce the amount of upload buffer space. | Follow the steps in To remove an Amazon EBS volume from your Amazon EC2–hosted gateway (p. 161). |

| Issue | Action to Take |
| --- | --- |
| Throughput to or from your Amazon EC2 gateway drops to zero. | • Verify the gateway instance is running. If the instance is starting due to a reboot, for example, wait for the instance to restart.<br>• Verify that the gateway IP has not changed. If the instance was stopped and then restarted, the IP address of the instance might have changed. In this case, you need to activate a new gateway.<br><br>You can view the throughput to and from your gateway from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway to AWS, see Measuring Performance Between Your Gateway and AWS (p. 188). |

**Warning**
For Amazon storage gateways that are deployed on an Amazon EC2 instance, accessing gateways over a public connection to the Internet is not supported. The elastic IP address of the Amazon EC2 instance cannot be used as the target address.

## Enabling AWS Support to Access a Gateway Hosted on an Amazon EC2 Instance

AWS Storage Gateway provides a local console you can use to perform several maintenance tasks, including enabling AWS Support to access your gateway to assist you with troubleshooting gateway issues. By default, AWS Support's access to your gateway is disabled. You enable this access through the Amazon EC2 local console. You log in to the Amazon EC2 local console via Secure Shell (SSH). To successfully log in via SSH, your instance's security group must have a rule allowing access via TCP port 22.

**Note**
If you add a new rule to an existing security group, the new rule applies to all instances that use that security group. For more information about security groups and how to add a security group rule, go to Amazon EC2 Security Groups in the *Amazon EC2 User Guide*.

To let AWS Support connect to your gateway, you first log in to the local console for the Amazon EC2 instance, navigate to the storage gateway's console, and then provide the access.

### To enable AWS support access to a gateway deployed on an Amazon EC2 instance

1. Log in to the local console for your Amazon EC2 instance. For instructions, go to Connect to Your Instancein the *Amazon EC2 User Guide*.

   You can use the following command to log in to the EC2 instance's local console.

   ```
   ssh –i PRIVATE-KEY sguser@INSTANCE-PUBLIC-DNS-NAME
   ```

   **Note**
   The *PRIVATE-KEY* is the `.pem` file containing the private certificate of the EC2 key pair that you used to launch the Amazon EC2 instance. For more information, go to Retrieving the Public Key for Your Key Pair in the *Amazon EC2 User Guide*.
   The *INSTANCE-PUBLIC-DNS-NAME* is the public Domain Name System (DNS) name of your Amazon EC2 instance that your gateway is running on. You obtain this public DNS name by selecting the Amazon EC2 instance in the EC2 console and clicking the **Description** tab.

   The local console looks like the following.

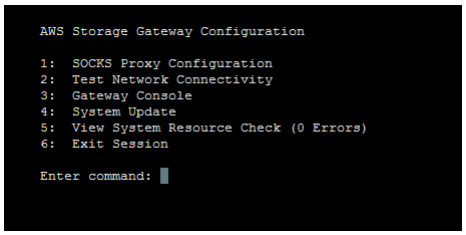


2. At the prompt, type **3** to open the AWS Storage Gateway console.

3. Type **h** to open the **AVAILABLE COMMANDS** window.
4. In the **AVAILABLE COMMANDS** window, type **`open-support-channel`** to connect to customer support for AWS Storage Gateway.
5. At the prompt, type a support service number for your gateway to connect to the AWS Storage Gateway support server. Choose any number between 2000 and 65000.

   Once the support channel is established, provide the support service number you selected to AWS Support so they can provide troubleshooting assistance.

6. When the support session is completed, type **q** to end it.
7. Type **`exit`** to exit the AWS Storage Gateway console.
8. Follow the console menus to log out of the AWS Storage Gateway instance.

## Troubleshooting Storage Volume Issues

The following table lists the most typical issues you might encounter when working with storage volumes.

| Issue | Action to Take |
| --- | --- |
| The AWS Storage Gateway console indicates that your volume has a status of UPLOAD BUFFER NOT CONFIGURED (p. 108). | Add upload buffer capacity to your gateway. You cannot use a gateway to store your application data if the upload buffer for the gateway is not configured. For more information, see To configure a local disk as an upload buffer for your gateway (p. 184). |
| The AWS Storage Gateway console indicates that your volume has a status of IRRECOVERABLE (p. 108). | You can no longer use this storage volume. You can try to delete the volume in the AWS Storage Gateway console. If there is data on the volume, then you can recover the data when you create a new storage volume based on the local disk of the VM that was initially used to create the storage volume. When you create the new storage volume, select **Preserve existing data**. For more information on this process, see Managing Storage Volumes in AWS Storage Gateway (p. 97). Delete pending snapshots of the volume before deleting the storage volume. For more information, see Deleting a Snapshot (p. 212).<br><br>If deleting the volume in the AWS Storage Gateway console does not work, then the disk allocated for the storage volume might have been improperly removed from the VM and cannot be removed from the appliance. |

| Issue | Action to Take |
| --- | --- |
| The AWS Storage Gateway console indicates that your volume has a status of PASS THROUGH (p. 108). | A volume can have PASS THROUGH (p. 108) status for several reasons. Some reasons require action, and some do not.<br><br>An example of where you should take action when your storage volume has the PASS THROUGH status is when your gateway has run out of upload buffer space. To verify if your upload buffer was exceeded in the past, you can view the `UploadBufferPercentUsed` metric in the Amazon CloudWatch console (see Monitoring the Upload Buffer (p. 191)). If your gateway has the PASS THROUGH status because it has run out of upload buffer space, you should allocate more upload buffer space to your gateway. Adding more buffer space will cause your storage volume to transition from PASS THROUGH to BOOTSTRAPPING (p. 108) to AVAILABLE (p. 108) automatically. While the volume has the BOOTSTRAPPING status, the gateway reads data off the volume's disk, uploads this data to Amazon S3, and catches up as needed. When the gateway has caught up and saved the storage volume data to Amazon S3, the volume status becomes AVAILABLE and snapshots can be started again. Note that when your storage volume has the PASS THROUGH or BOOTSTRAPPING status, you can continue to read and write data from the storage volume disk. For more information about adding more upload buffer space, see Managing the Upload Buffer (Gateway-Stored) (p. 178).<br><br>To take action before the upload buffer is exceeded, you can set a threshold alarm on a gateway's upload buffer. For more information, see To set an upper threshold alarm for a gateway's upload buffer (p. 192).<br><br>Another example of not needing to take action when a storage volume has the PASS THROUGH status is when the storage volume is waiting to be bootstrapped because another volume is currently being bootstrapped. The gateway bootstraps volumes one at a time.<br><br>Infrequently, the PASS THROUGH status can indicate that a disk allocated for an upload buffer has failed. In this is the case, you should remove the disk. For more information, see Removing Upload Buffer Capacity (p. 173) or Removing Upload Buffer Capacity (Gateway-Stored) (p. 181). |
| You want to verify the integrity of your volumes and fix possible errors on the volumes. | For gateways that use Microsoft Windows initiators to connect to volumes, you can use the Windows CHKDSK utility to verify the integrity of your volumes and fix errors on the volumes. Windows can automatically run the CHKDSK tool when volume corruption is detected, or you can run it yourself. |
| Your storage volume's iSCSI target does not show up in the Disk Management Console in Windows. | Check that you have configured the upload buffer for the gateway. For more information, see To configure a local disk as an upload buffer for your gateway (p. 184). |
| You want to change the iSCSI target name of your storage volume. | You cannot change the iSCSI target name of your storage volume except by deleting the volume and adding it again with a new target name. If you do so, you can preserve the data on the volume. For information about creating a storage volume, see Managing Storage Volumes (Gateway-Cached) (p. 97) or Managing Storage Volumes (Gateway-Stored) (p. 100). |

| Issue | Action to Take |
| --- | --- |
| Your scheduled snapshot of a storage volume did not occur. | Check if your volume has the PASS THROUGH (p. 108) status, or if the gateway's upload buffer was filled just prior to the scheduled snapshot time. You can check the `UploadBufferPercentUsed` metric for the gateway in the Amazon CloudWatch console and create an alarm for this metric. For more information, see Monitoring the Upload Buffer (p. 191) and To set an upper threshold alarm for a gateway's upload buffer (p. 192). |
| You need to remove a storage volume because it isn't needed, or you need to replace a storage volume disk that has failed. | You should remove the volume first using the AWS Storage Gateway console (see To remove a storage volume (p. 100)) and then use the hypervisor client to remove the backing storage.<br>• For VMware ESXi, remove the backing storage as described in To remove the underlying local disk (VMware ESXi) (p. 103).<br>• For Microsoft Hyper-V, remove the backing storage as described in To remove the underlying local disk (Microsoft Hyper-V) (p. 104). |
| Throughput from your application to a storage volume has dropped to zero. | • Check that your storage volume's **Host IP** address matches one of the addresses that appears in the vSphere client on the **Summary** tab. You can find the **Host IP** address for a storage volume in the AWS Storage Gateway console in the **ISCSI Target Info** tab for the storage volume. A discrepancy in the IP address can occur, for example, when you assign a new static IP address to your gateway. If there is a discrepancy, restart your gateway from the AWS Storage Gateway console as shown in Shutting Down and Starting a Gateway Using the AWS Storage Gateway Console (p. 127). After the restart, the **Host IP** address in the **ISCSI Target Info** tab for a storage volume should match an IP address shown in the vSphere client on the **Summary** tab for the gateway.<br>• Check to see if **IPAddressNotFound** appears in the **Host IP** box for the storage volume. For example, this message can appear when you create a storage volume associated with an IP address of a network adapter of a gateway with two or more network adapters. When you remove or disable the network adapter that the storage volume is associated with, the **IPAddressNotFound** message is displayed. To address this issue, delete the storage volume and then re-create it preserving its existing data. For more information, see Managing Storage Volumes in AWS Storage Gateway (p. 97).<br>• Check that the iSCSI initiator your application uses is correctly mapped to the iSCSI target for the storage volume. For more information about connecting to storage volumes, see Connecting to Storage Volumes (p. 111).<br><br>You can view the throughput for storage volumes and create alarms from the Amazon CloudWatch console. For more information about measuring throughput from your application to a storage volume, see Measuring Performance Between Your Application and Gateway (p. 187). |

| Issue | Action to Take |
| --- | --- |
| A cache disk in your gateway encounters a failure. | If one or more cache disks in your gateway encounters a failure, the gateway prevents read and write operations to your virtual tapes and volumes. To resume normal functionality, reconfigure your gateway as described following:<br>• If the cache disk is inaccessible or unusable, delete the disk from your gateway configuration.<br>• If the cache disk is still accessible and useable, reconnect it to your gateway.<br>**Note**<br>If you delete a cache disk, tapes or volumes that have clean data (that is, for which data in the cache disk and Amazon S3 are synchronized) will continue to be available when the gateway resumes normal functionality. For example, if your gateway has three cache disks and you delete two, tapes or volumes that are clean will have AVAILABLE status. Other tapes and volumes will have IRRECOVERABLE status.<br>If you use ephemeral disks as cache disks for your gateway or mount your cache disks on an ephemeral drive, your cache disks will be lost when you shut down the gateway. Shutting down the gateway when your cache disk and Amazon S3 are not synchronized can result in data loss. As a result, we don't recommend using ephemeral drives or disks. |

## Troubleshooting Snapshot Issues

The following table lists issues you might encounter when working with snapshots.

| Issue | Action to Take |
| --- | --- |
| A storage volume snapshot remains in PENDING state longer than expected. | If the gateway VM crashes unexpectedly or the status of a volume changes to PASS THROUGH (p. 108) or IRRECOVERABLE (p. 108) before a volume snapshot completes, then the status of the snapshot will remain in PENDING status and the snapshot does not successfully complete. If this is the case, we recommend that you delete the snapshot. For more information, see Deleting a Snapshot (p. 212). Once the volume returns to AVAILABLE (p. 108) status, create a new snapshot of the volume. |

## Using Recovery Snapshots (Gateway-Cached)

AWS Storage Gateway provides recovery points for each volume in a gateway-cached volume architecture. A volume recovery point is a point in time at which all data of the volume is consistent and from which you can create a snapshot. You can use the snapshot to create a new volume in the event that your gateway becomes unreachable or one gateway-cached volume becomes irrecoverable.

Volume recovery points are maintained automatically for each gateway-cached volume. You can also take snapshots on an ad hoc basis or set up a snapshot schedule for the volume. For more information about snapshots, see Working with Snapshots (p. 201).

When the gateway becomes unreachable (such as when you shut it down), you have the option of creating a snapshot from a volume recovery point.

### To create and use a recovery snapshot of a volume from an unreachable gateway

1. On the **AWS Storage Gateway** console navigation pane, select the unreachable gateway.

   A list of volumes of the gateway is displayed.

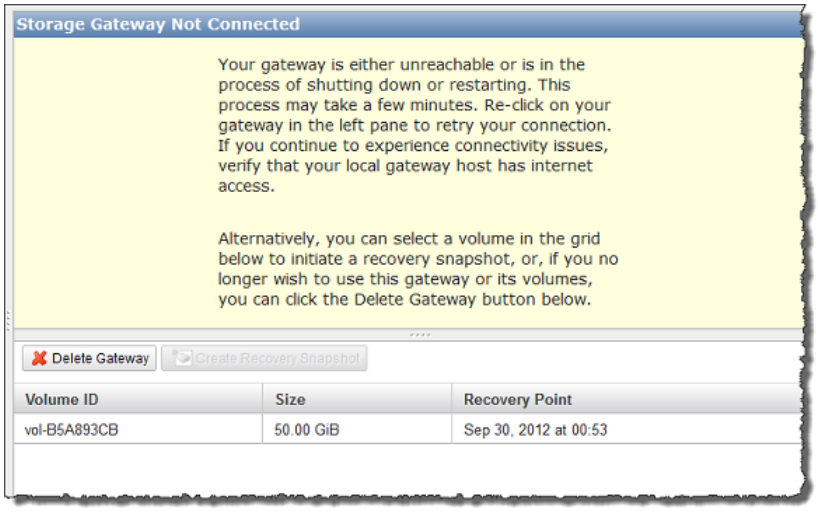
Storage Gateway Not Connected
Your gateway is either unreachable or is in the process of shutting down or restarting. This process may take a few minutes. Re-click on your gateway in the left pane to retry your connection. If you continue to experience connectivity issues, verify that your local gateway host has internet access.
Alternatively, you can select a volume in the grid below to initiate a recovery snapshot, or, if you no longer wish to use this gateway or its volumes, you can click the Delete Gateway button below.
Delete Gateway
Create Recovery Snapshot
Volume ID
Size
Recovery Point
vol-B5A893CB
50.00 GiB
Sep 30, 2012 at 00:53

2. Select the volume from which to create a recovery snapshot.
3. Click the **Create Recovery Snapshot** button.

   AWS Storage Gateway initiates the snapshot process.

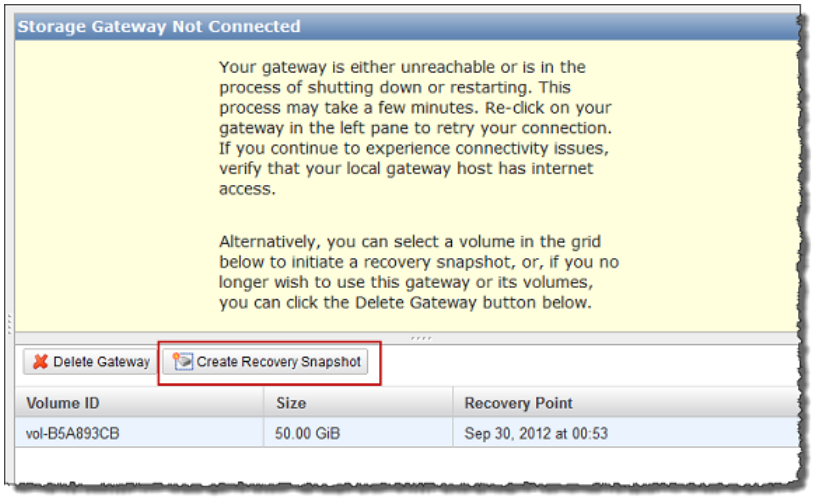
Storage Gateway Not Connected
Your gateway is either unreachable or is in the process of shutting down or restarting. This process may take a few minutes. Re-click on your gateway in the left pane to retry your connection. If you continue to experience connectivity issues, verify that your local gateway host has internet access.
Alternatively, you can select a volume in the grid below to initiate a recovery snapshot, or, if you no longer wish to use this gateway or its volumes, you can click the Delete Gateway button below.
Delete Gateway
Create Recovery Snapshot
Volume ID
Size
Recovery Point
vol-B5A893CB
50.00 GiB
Sep 30, 2012 at 00:53

4. In the **Recovery Point Created** dialog box, click **Close**.

5. Find the snapshot using the steps in the procedure Finding a Snapshot (p. 203).
6. Restore the snapshot using one of the procedures in Restoring a Snapshot (p. 222).

# Gateway–Virtual Tape Library (Gateway-VTL)

Gateway-virtual tape library (VTL) offers a durable, cost-effective solution to offline data archiving. The VTL interface lets you use your existing tape-based backup infrastructure to store data on virtual tape cartridges, which you create on your gateway-VTL. Each gateway-VTL is preconfigured with a media changer and tape drives, which are available to your existing client backup applications as Internet Small Computer System Interface (iSCSI) devices. You add virtual tape cartridges as you need to archive your data.

For more information about the gateway-VTL architecture, see Gateway–Virtual Tape Library (Gateway-VTL) Architecture (p. 5). For information about gateway concepts and storage solutions offered, see What Is AWS Storage Gateway? (p. 1)

The following sections provide information about setting up and managing a gateway-VTL. You can use the AWS Storage Gateway console to perform these tasks, or you can perform these tasks programmatically using the AWS Storage Gateway API. For more information on using this API, see API Reference for AWS Storage Gateway (p. 379) or the AWS SDK libraries.

**Topics**

- Setting Up Your Gateway-VTL (p. 240)
- Managing Your Activated Gateway-VTL (p. 276)

## Setting Up Your Gateway-VTL

As you follow the steps in this section, you will use the **Setup and Activate Gateway** wizard in the AWS Storage Gateway console. At several steps in the wizard, you will perform tasks outside of the console and then return. If your session times out or the browser closes, you can return to the console.

**Topics**

- Step 1: Before You Begin (p. 241)
- Step 2: Configure Local Storage and Alarms (Gateway-VTL) (p. 241)
- Step 3: Create Virtual Tapes Using the AWS Storage Gateway Console (p. 243)
- Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245)
- Step 5: Test the Gateway-VTL Setup (p. 249)

- Step 6: Clean Up After the Example Exercise (p. 275)
- Where Do I Go from Here? (p. 276)

# Step 1: Before You Begin

Before you begin your gateway setup, you need to do the following:

1. Sign up for an AWS account. To use AWS Storage Gateway, you need an AWS account.
2. Deploy and activate a gateway. Steps to deploy and activate gateway are the same regardless of the type of storage solution you are deploying.

You can find instructions for deploying and activating a gateway in the Getting Started section that is common to all gateways, Getting Started with AWS Storage Gateway (p. 9). Before going on, you must first complete these steps. When you have done so, continue to the next step.

**Next Step**

Step 2: Configure Local Storage and Alarms (Gateway-VTL) (p. 241)

# Step 2: Configure Local Storage and Alarms (Gateway-VTL)

Earlier in this Getting Started exercise, you created two disks, one for an upload buffer and the other for cache storage. Now, you will configure your gateway to use these disks. Optionally, you can also create alarms to monitor your disks.

In a gateway-VTL setup, after your gateway is activated, the **Create Tapes** button is displayed. If activation is not successful, see Troubleshooting Your Gateway-VTL (p. 363) for possible solutions.

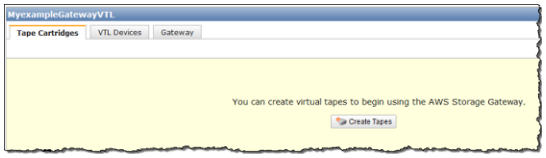


When you click **Create Tapes**, the wizard walks you through how to configure local storage, optionally create alarms, and create virtual tapes for your gateway.

> **Note**
> When creating virtual tapes for a gateway for the first time, you must first configure local storage for your gateway.

### To configure local storage (cache storage and upload buffer)

1. Click **Create Tapes** to start the **Configure Your Activated Gateway** wizard.
2. In the **Configure Your Activated Gateway** wizard, configure local working storage.

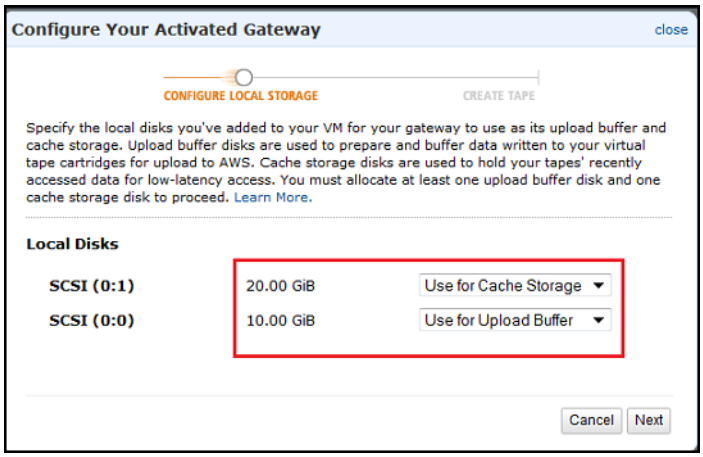
Configure Your Activated Gateway
close
CONFIGURE LOCAL STORAGE
CREATE TAPE
Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage. Upload buffer disks are used to prepare and buffer data written to your virtual tape cartridges for upload to AWS. Cache storage disks are used to hold your tapes' recently accessed data for low-latency access. You must allocate at least one upload buffer disk and one cache storage disk to proceed. Learn More.
Local Disks
SCSI (0:1)
20.00 GiB
Use for Cache Storage
SCSI (0:0)
10.00 GiB
Use for Upload Buffer
Cancel
Next

a. Of the two disks you provisioned, select one as the upload buffer and the other as cache storage, as shown in the preceding screenshot.

b. Click **Next**.

Optionally, you can create alarms to monitor the storage use of the two disks you created.

### To configure alarms

1. In the upload buffer alarm page, configure the alarm for upload buffer use.

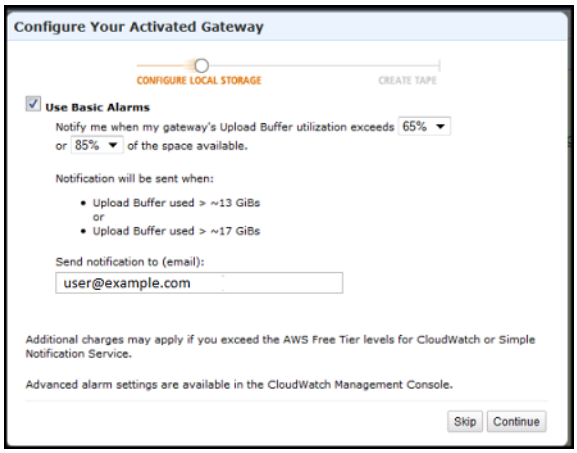
Configure Your Activated Gateway
CONFIGURE LOCAL STORAGE
CREATE TAPE
Use Basic Alarms
Notify me when my gateway's Upload Buffer utilization exceeds 65%
or 85% of the space available.
Notification will be sent when:
Upload Buffer used > ~13 GiBs
or
Upload Buffer used > ~17 GiBs
Send notification to (email):
user@example.com
Additional charges may apply if you exceed the AWS Free Tier levels for CloudWatch or Simple Notification Service.
Advanced alarm settings are available in the CloudWatch Management Console.
Skip
Continue

   a. Using the two drop-down lists, select utilization percentages to set two upload buffer alarms.

      For example, you can set the first threshold (the lower percentage value) to a percentage use of the upload buffer that, if exceeded, you want to be warned about. You can set the second threshold (the higher percentage value) to a percentage use of the upload buffer that, if exceeded, is cause for action. The action might be adding more upload buffer space.

      After you complete this step, you can go to the CloudWatch console at any time and change the alarm thresholds.

   b. Type an email address to which alarm notifications should be sent.

   c. Click **Continue**.

Two alarms are created. You can name the alarms—for example, *MyGatewayBufferUtilizationAlarm*.

d. Check the email address you provided for a subscription confirmation email, and follow the instructions in that email to confirm your subscription to the related Amazon Simple Notification Service (Amazon SNS) topic. After you have confirmed your subscription, you will receive an email when either threshold you specified is exceeded.

2. Click **Continue** to create the alarms.
3. In the cached storage alarms page, configure the alarms to notify you about cache storage use.

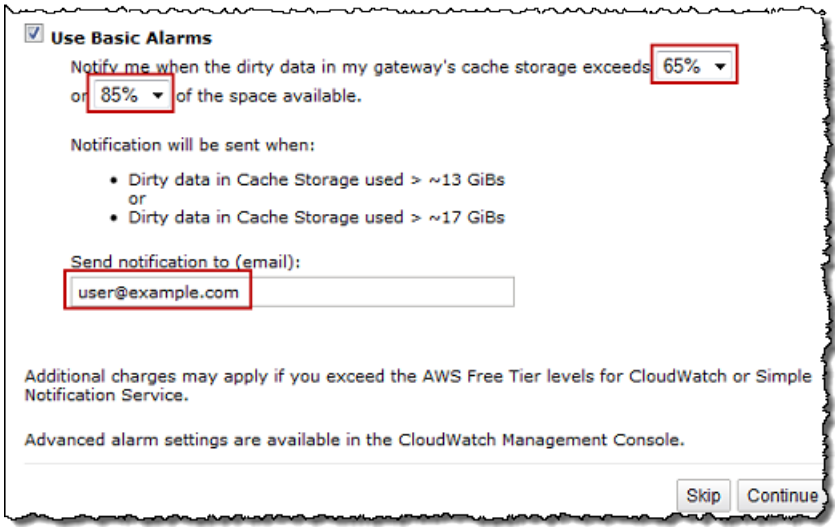
Use Basic Alarms
Notify me when the dirty data in my gateway's cache storage exceeds 65%
or 85% of the space available.
Notification will be sent when:
Dirty data in Cache Storage used > ~13 GiBs
or
Dirty data in Cache Storage used > ~17 GiBs
Send notification to (email):
user@example.com
Additional charges may apply if you exceed the AWS Free Tier levels for CloudWatch or Simple Notification Service.
Advanced alarm settings are available in the CloudWatch Management Console.
Skip
Continue

Follow the instructions in step 1 of the To configure alarms (p. 242) procedure to configure the alarms.

You can name the alarms—for example, *MyGatewayCacheBufferAlarm*.

4. Click **Continue** to create the alarm.

**Next Step**

Step 3: Create Virtual Tapes Using the AWS Storage Gateway Console (p. 243)

# Step 3: Create Virtual Tapes Using the AWS Storage Gateway Console

Now that you have activated your gateway-VTL and configured local storage, you will create virtual tapes.

**Note**
You are charged only for the amount of data you write to the tape, not the tape capacity.

The following screenshot shows the **Create Virtual Tapes** page.

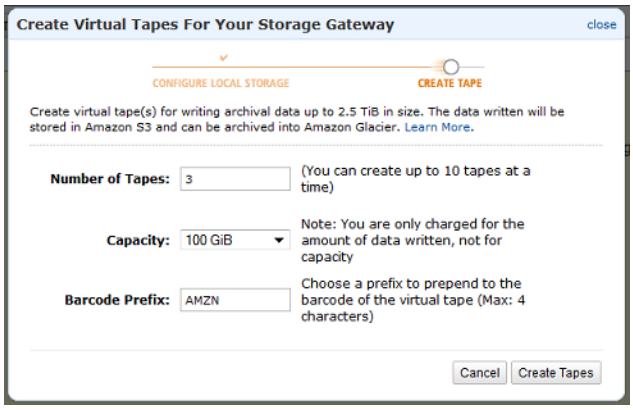
Create Virtual Tapes For Your Storage Gateway
close
CONFIGURE LOCAL STORAGE
CREATE TAPE
Create virtual tape(s) for writing archival data up to 2.5 TiB in size. The data written will be stored in Amazon S3 and can be archived into Amazon Glacier. Learn More.
Number of Tapes:
3
(You can create up to 10 tapes at a time)
Capacity:
100 GiB
Note: You are only charged for the amount of data written, not for capacity
Barcode Prefix:
AMZN
Choose a prefix to prepend to the barcode of the virtual tape (Max: 4 characters)
Cancel
Create Tapes

### To create virtual tapes

1. In the **Number of Tapes** text box, type the number of tapes you want to create. For this exercise, you can create three tapes. You can create a maximum of 1500 tapes per gateway. However, you can create a maximum of 10 virtual tapes at a time. For more information, see Appendix F: AWS Storage Gateway Limits (p. 421).
2. In the **Capacity** text box, select the size of virtual tapes you want to create. For this exercise, you can enter 100 GiB. For information about capacity limits, see Appendix F: AWS Storage Gateway Limits (p. 421).
3. In the **Barcode Prefix** text box, type the prefix you want to prepend to the barcode of your virtual tapes. For this exercise, you can accept the default.

   **Note**

   Virtual tapes are uniquely identified by a barcode. You can add a prefix to the barcode. The prefix is optional, but you can use it for your virtual tapes. The prefix must be uppercase letters (A–Z) and must be 1 to 4 characters long.
4. Click **Create Tapes**.

   The confirmation page shows the iSCSI targets on your gateway. You can view this information in the AWS Storage Gateway console at any time by selecting your gateway and navigating to the **VTL Devices** tab.

   The following screenshot shows the confirmation page for the devices that were created for your gateway.

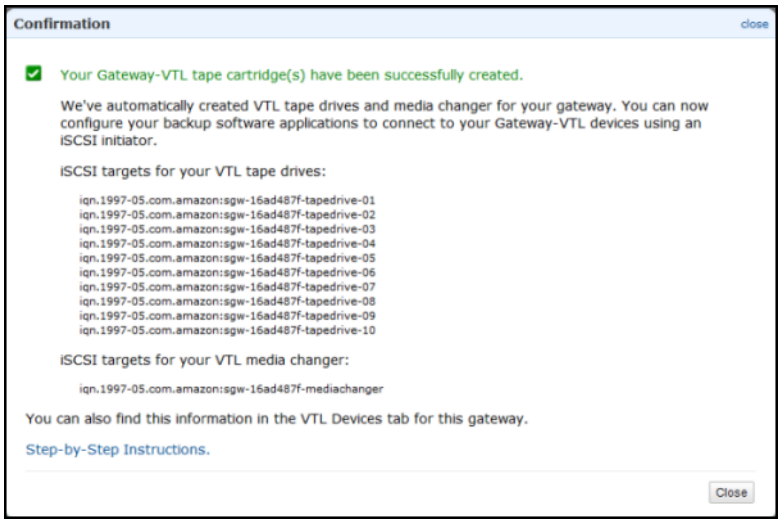
Confirmation
close
Your Gateway-VTL tape cartridge(s) have been successfully created.
We've automatically created VTL tape drives and media changer for your gateway. You can now configure your backup software applications to connect to your Gateway-VTL devices using an iSCSI initiator.
iSCSI targets for your VTL tape drives:
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-01
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-02
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-03
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-04
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-05
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-06
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-07
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-08
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-09
iqn.1997-05.com.amazon:sgw-16ad487f-tapedrive-10
iSCSI targets for your VTL media changer:
iqn.1997-05.com.amazon:sgw-16ad487f-mediachanger
You can also find this information in the VTL Devices tab for this gateway.
Step-by-Step Instructions.
Close

5. Click **Close** to close the **Confirmation** page.

The virtual tapes you have created appear in the AWS Storage Gateway console. The console shows the barcode, capacity, and status of the virtual tapes. The status of the virtual tapes is initially set to **CREATING** when the virtual tapes are being created. After the tapes are created, their status changes to **AVAILABLE**. For more information, see Managing Tapes in the VTL (p. 277).

**Next Step**

Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245)

# Step 4: Connect Your Gateway-VTL Devices to Your Windows Client

In this Getting Started exercise, you use the Windows iSCSI initiator to connect to your gateway's VTL devices. At the end of this procedure, the devices become available as local devices on your Windows client. After this step, you configure your backup software to access the devices. For instructions on accessing the iSCSI VTL devices from Linux, see Connecting from a Red Hat Client to Your VTL Devices (p. 291).

### To connect your Windows client to the VTL devices

1. On the **Start** menu of your Windows client computer, type **`iscsicpl.exe`** in the **Search Programs and files** box, locate the iSCSI initiator program and then run it.

   **Note**
   You must have administrator rights on the client computer to run the iSCSI initiator.

2. If prompted, click **Yes** to start the Microsoft iSCSI initiator service.

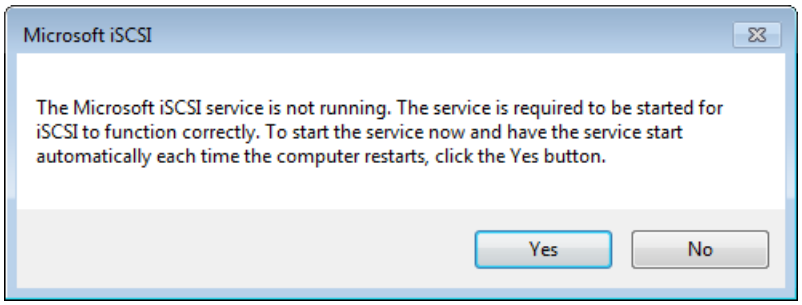
Microsoft iSCSI
The Microsoft iSCSI service is not running. The service is required to be started for iSCSI to function correctly. To start the service now and have the service start automatically each time the computer restarts, click the Yes button.
Yes
No

3. In the **iSCSI Initiator Properties** dialog box, click the **Discovery** tab, and then click the **Discover Portal** button.

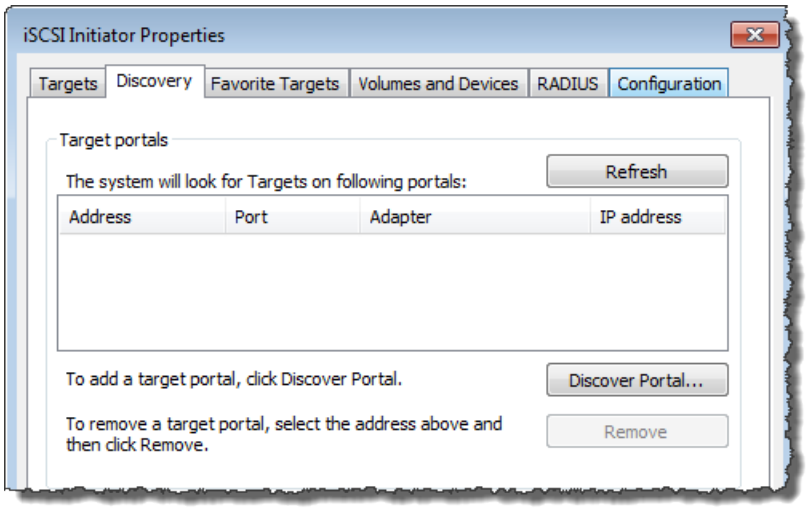


4. In the **IP address or DNS name** box of the **Discover Target Portal** dialog box, type the IP address of your gateway-VTL, and then click **OK**. You can obtain the IP address of your gateway in the **Gateway** tab on the AWS Storage Gateway console. If you deployed your gateway on an Amazon EC2 instance, you can obtain the public IP address or public DNS in the **Description** tab on the Amazon EC2 console.

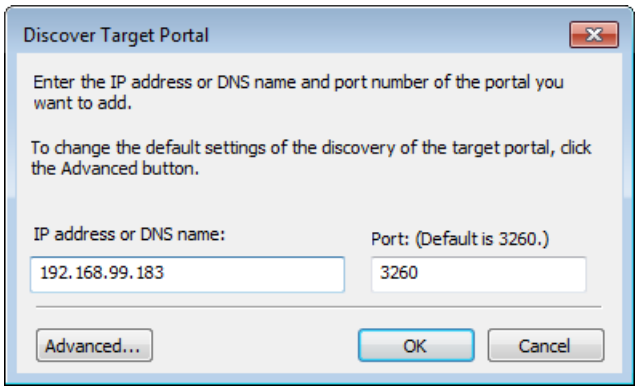
Discover Target Portal
Enter the IP address or DNS name and port number of the portal you want to add.
To change the default settings of the discovery of the target portal, click the Advanced button.
IP address or DNS name:
Port: (Default is 3260.)
192.168.99.183
3260
Advanced...
OK
Cancel

5. Select the **Targets** tab and click **Refresh**. All 10 tape drives and the medium changer appear in the **Discovered targets** box. The status for the targets is **Inactive**.

   The following screenshot shows the discovered targets.

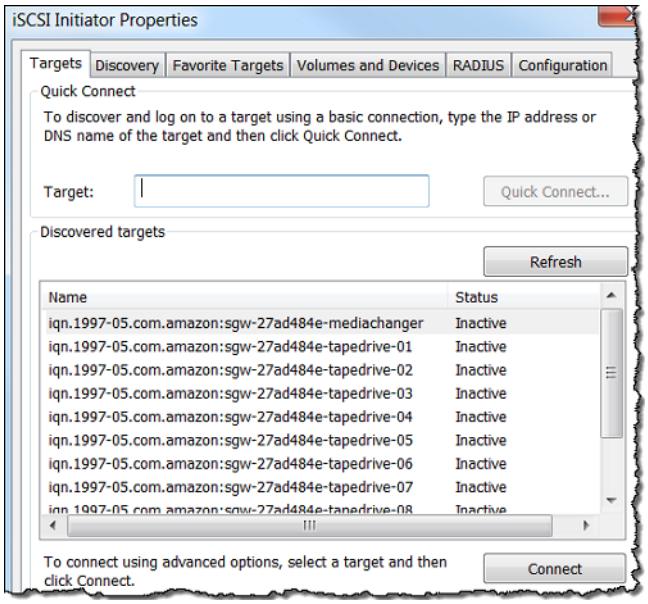


6. Select the first device and click **Connect**. You connect the devices one at a time.
7. In the **Connect to Target** dialog box, click **OK**.
8. Repeat steps 6 and 7 for each of the devices to connect all of them, and then click **OK** in the **iSCSI Initiator Properties** dialog box.
9. On a Windows client, the driver provider for the tape drive must be Microsoft. Use the following procedure to verify the driver provider, and update the driver and provider if necessary.

    1. On your Windows client, start Device Manager.
    2. Expand **Tape drives**, right-click on a tape drive, and select **Properties**.

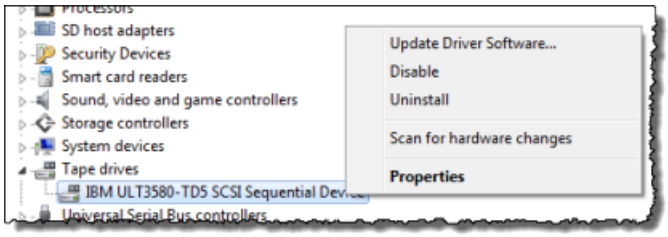


    3. In the **Driver** tab of the **Device Properties** dialog box, verify the **Driver Provider** is Microsoft.

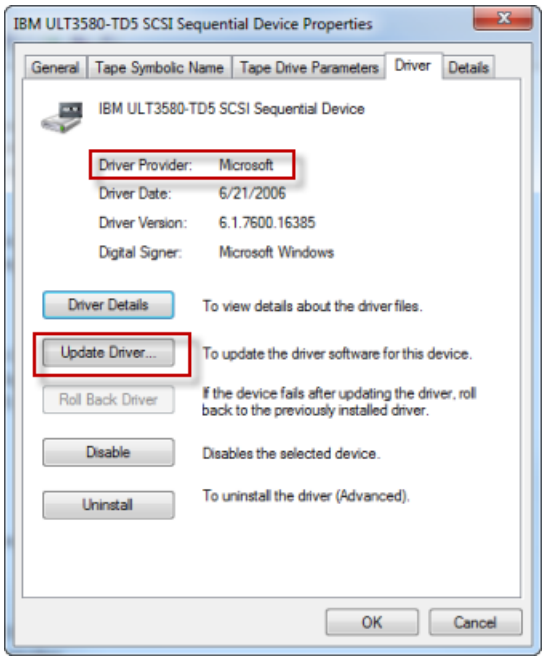


4. If the **Driver Provider** is not Microsoft, set the value as follows:

    1. Click **Update Driver**.
    2. In the **Update Driver Software** dialog box, click **Browse my computer for driver software**.

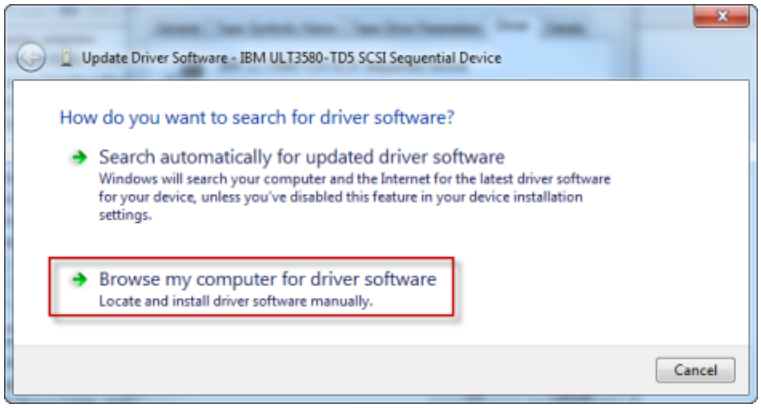
Update Driver Software - IBM ULT3580-TD5 SCSI Sequential Device
How do you want to search for driver software?
Search automatically for updated driver software
Windows will search your computer and the Internet for the latest driver software for your device, unless you've disabled this feature in your device installation settings.
Browse my computer for driver software
Locate and install driver software manually.
Cancel

    3. In the **Update Driver Software** dialog box, click **Let me pick from a list of device drivers on my computer**.

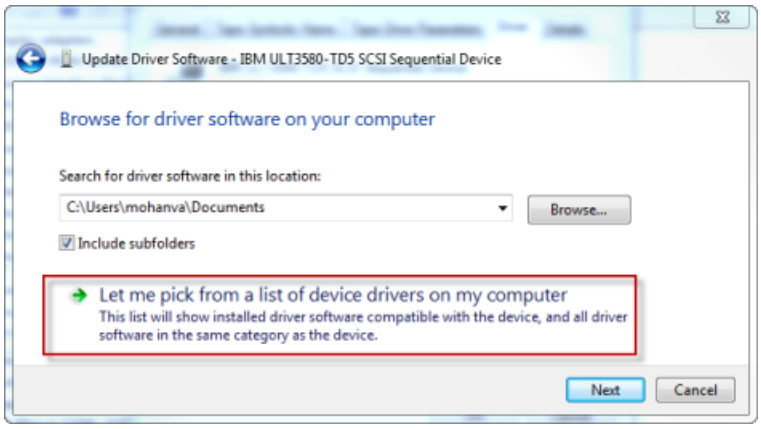
Update Driver Software - IBM ULT3580-TD5 SCSI Sequential Device
Browse for driver software on your computer
Search for driver software in this location:
C:\Users\mohanva\Documents
Browse...
Include subfolders
Let me pick from a list of device drivers on my computer
This list will show installed driver software compatible with the device, and all driver software in the same category as the device.
Next
Cancel

4. Select **LTO Tape drive** and click **Next**.

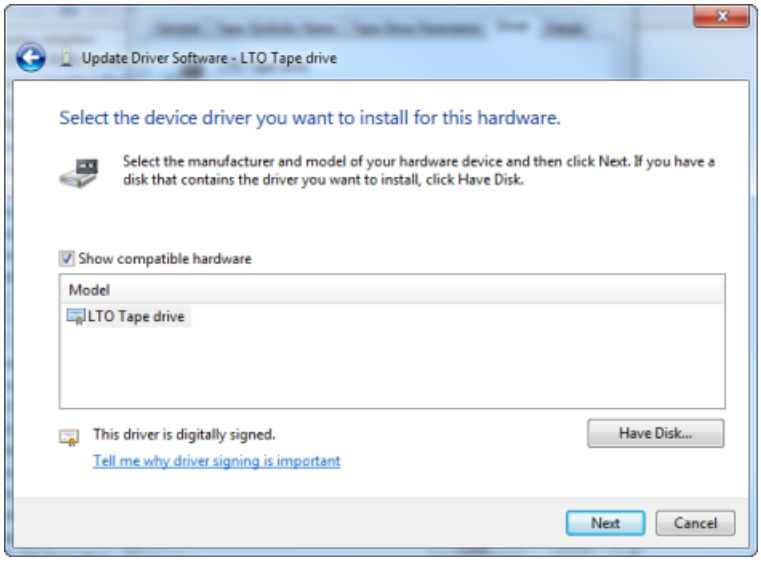
Update Driver Software - LTO Tape drive
Select the device driver you want to install for this hardware.
Select the manufacturer and model of your hardware device and then click Next. If you have a disk that contains the driver you want to install, click Have Disk.
Show compatible hardware
Model
LTO Tape drive
This driver is digitally signed.
Tell me why driver signing is important
Have Disk...
Next
Cancel

5. Click **Close** to close the **Update Driver Software** window, and verify that the **Driver Provider** value is now set to Microsoft.
6. Repeat steps 9.2 through 9.5 to update all the tape drives.

**Next Step**

Step 5: Test the Gateway-VTL Setup (p. 249)

# Step 5: Test the Gateway-VTL Setup

You test your gateway-VTL setup by performing the following tasks using your backup software:

1. Configure the backup software to detect your storage devices.
2. Back up data to a tape.
3. Archive the tape.
4. Retrieve the tape from the archive.

5. Restore data from the tape.

You can test your setup with one of the following types of backup software:

- Testing Your Setup by Using Symantec NetBackup Version 7.x (p. 250).
- Testing Your Setup by Using Symantec Backup Exec (p. 263).
- Testing Your Setup by Using Microsoft System Center 2012 R2 Data Protection Manager (p. 267).
- Testing Your Setup by Using Veeam Backup & Replication (p. 270).
- Testing Your Setup by Using Dell NetVault Backup (p. 273).

# Testing Your Setup by Using Symantec NetBackup Version 7.x

AWS Storage Gateway-VTL supports backing up data to virtual tapes and archiving the tapes in a virtual tape shelf (VTS) by using Symantec NetBackup version 7.x. In this topic, we provide documentation on how to configure the NetBackup software for a gateway–virtual tape library (VTL) and perform a backup. For detailed information about how to use NetBackup, go to the NetBackup documentation on the Symantec website. For Symantec support information on hardware compatibility, go to the Symantec NetBackup Enterprise Server and Server 7.x Hardware Compatibility List on the Symantec website.

This topic shows you to configure storage, write data to a tape, archive a tape in the VTS, and restore the data.

**Topics**

- Configuring NetBackup Storage Devices (p. 250)
- Backing Up Sample Data to a Tape on Your Gateway-VTL (p. 253)
- Archiving the Tape (p. 259)
- Retrieving the Archived Tape from the VTS Back to Your Gateway-VTL (p. 260)
- Restoring Data from the Tape (p. 261)

## Configuring NetBackup Storage Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Symantec NetBackup version 7.x storage to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245).

### To configure NetBackup to use storage devices on your gateway-VTL

1. Open the NetBackup Administration Console and run it as an administrator.

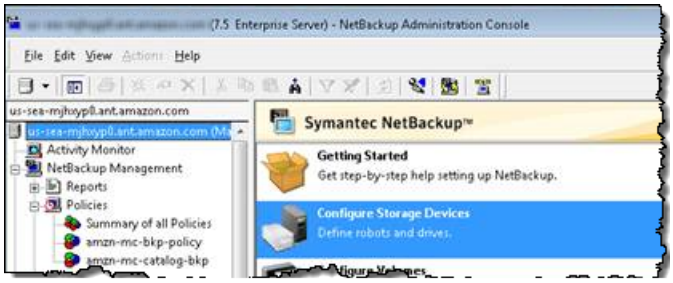


2. Click **Configure Storage Devices** to open the Device Configuration wizard.
3. Click **Next**. The NetBackup software detects your computer as a device host.

4. In the **Device Hosts** column, select your computer, and then click **Next**. The NetBackup software scans your computer for devices and discovers all devices.

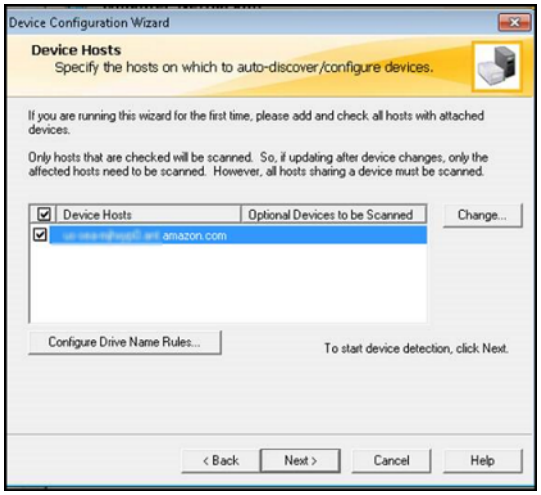
Device Configuration Wizard
Device Hosts
Specify the hosts on which to auto-discover/configure devices.
If you are running this wizard for the first time, please add and check all hosts with attached devices.
Only hosts that are checked will be scanned. So, if updating after device changes, only the affected hosts need to be scanned. However, all hosts sharing a device must be scanned.
Device Hosts
Optional Devices to be Scanned
Change...
amazon.com
Configure Drive Name Rules...
To start device detection, click Next.
< Back
Next >
Cancel
Help

5. In the **Scanning Hosts** page, click **Next**, and then click **Next**. The NetBackup software finds all 10 tape drives and the medium changer on your computer.

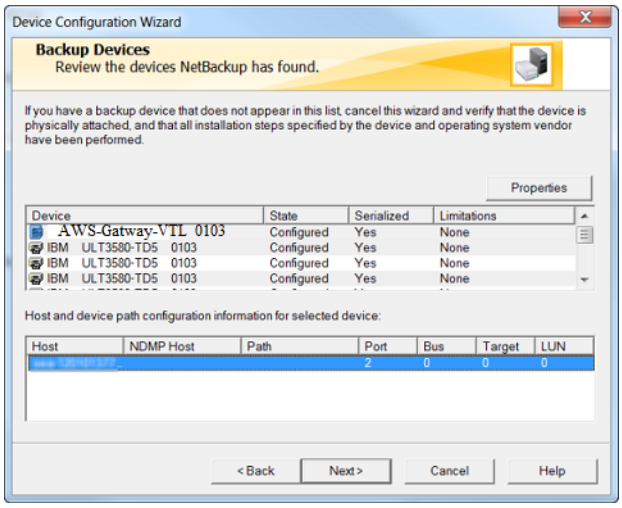


6. In the **Backup Devices** window, click **Next**.
7. In the **Drag and Drop Configuration** window, verify that your medium changer is selected, and then click **Next.**
8. In the dialog box that appears, click **Yes** to save the configuration on your computer. The NetBackup software updates the device configuration.
9. When the update is completed, click **Next** to make the devices available to the NetBackup software.
10. In the **Finished!** window, click **Finish**.

### To verify your devices in the NetBackup software

1. In the NetBackup Administration Console, expand the **Media and Device Management** node, and then expand the **Devices** node. Click **Drives** to display all the tape drives.

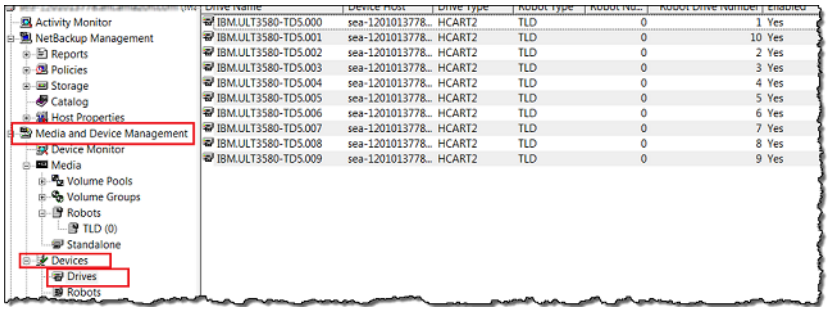


2. In the **Devices** node, click **Robots** to display all your medium changers. In the NetBackup software, the medium changer is called a *robot*.
3. In the **All Robots** pane, right-click **TLD(0)** (that is, your robot), and then click **Inventory Robot**.
4. In the **Robot Inventory** window, verify that your host is selected from the **Device-Host** list located in the **Select robot** category.
5. Verify that your robot is selected from the **Robot** list.
6. In the **Robot Inventory** window, select **Update volume configuration**, select **Preview changes**, select **Empty media access port prior to update**, and then click **Start**.

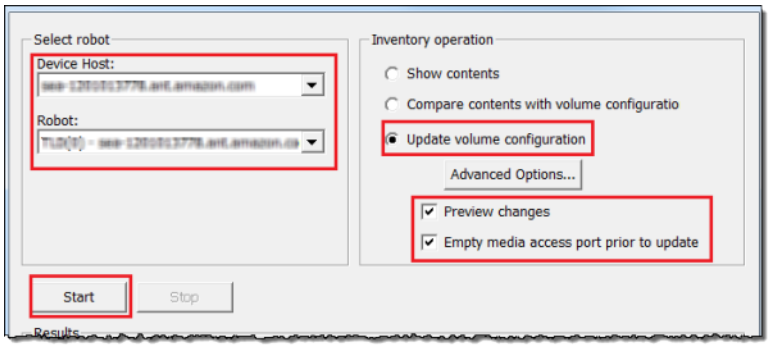


The process then inventories your medium changer and virtual tapes in the NetBackup Enterprise Media Management (EMM) database. NetBackup stores media information, device configuration, and tape status in the EMM.

7. In the **Robot Inventory** window, click **Yes** once the inventory is complete. Clicking **Yes** here updates the configuration and moves virtual tapes found in import/export slots to the virtual tape library.

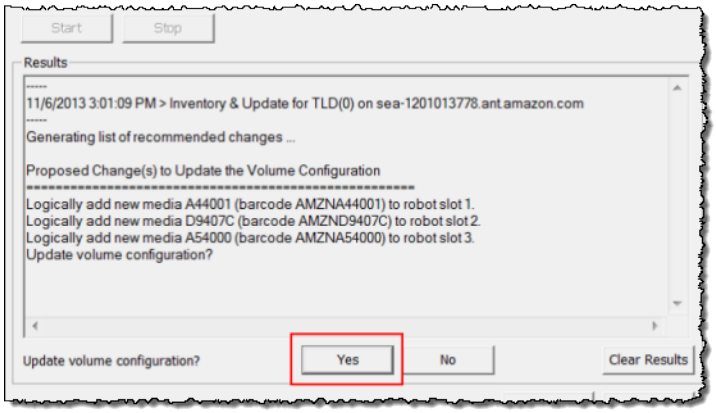
Start
Stop
Results
11/6/2013 3:01:09 PM > Inventory & Update for TLD(0) on sea-1201013778.ant.amazon.com
Generating list of recommended changes ...
Proposed Change(s) to Update the Volume Configuration
Logically add new media A44001 (barcode AMZNA44001) to robot slot 1.
Logically add new media D9407C (barcode AMZND9407C) to robot slot 2.
Logically add new media A54000 (barcode AMZNA54000) to robot slot 3.
Update volume configuration?
Update volume configuration?
Yes
No
Clear Results

For example, the following screenshot shows three virtual tapes found in the import/export slots.

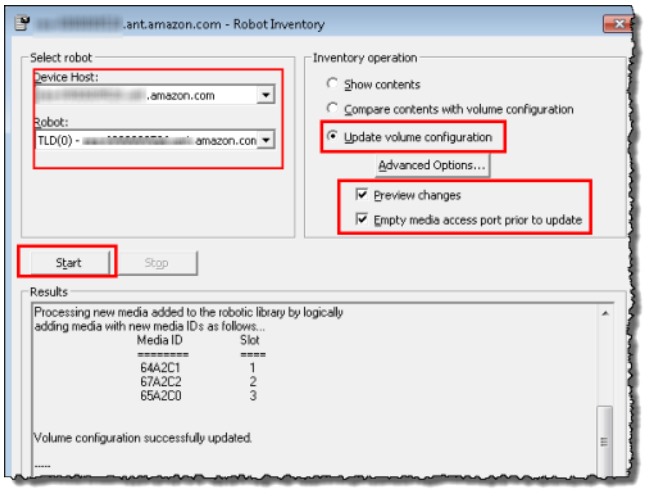


8. Close the **Robot Inventory** window.
9. In the **Media** node, expand the **Robots** node and click **TLD(0)** to show all virtual tapes that are available to your robot (medium changer).

   **Note**
   If you have previously connected other devices to the NetBackup software, you might have multiple robots. Make sure you select the right robot.

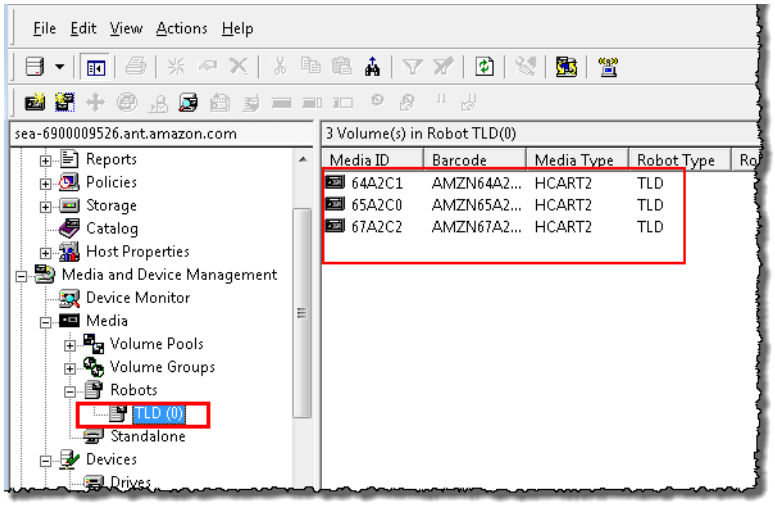


Now that you have connected your devices and made them available to your backup software, you are ready to test your gateway. To test your gateway, you back up data onto the virtual tapes you created and archive the tapes in the virtual tape shelf (VTS), which is backed by Amazon Glacier.

## Backing Up Sample Data to a Tape on Your Gateway-VTL

You test the gateway-VTL setup by backing up data onto your virtual tapes.

**Note**
You should back up only a small amount of data for this Getting Started exercise, because there are costs associated with storing, archiving, and retrieving data. For pricing information, go to Pricing on the AWS Storage Gateway detail page.

### To create a volume pool

A *volume pool* is a collection of virtual tapes to use for a backup.

1. Start the NetBackup Administration Console.
2. Expand the **Media** node, right-click **Volume Pool**, and then click **New**. The **New Volume Pool** dialog box appears.

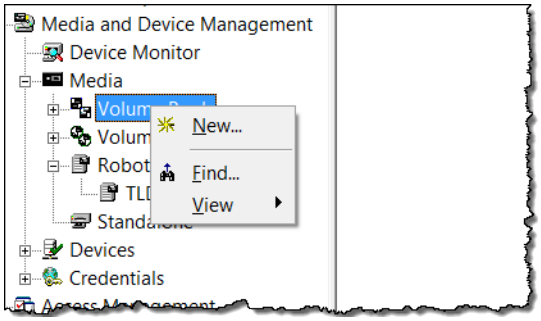


3. In the **Name** box, type a name for your volume pool.
4. In the **Description** box, type a description for the volume pool, and then click **OK**. The volume pool you just created is added to the volume pool list.

   The following screenshot shows a list of volume pools.

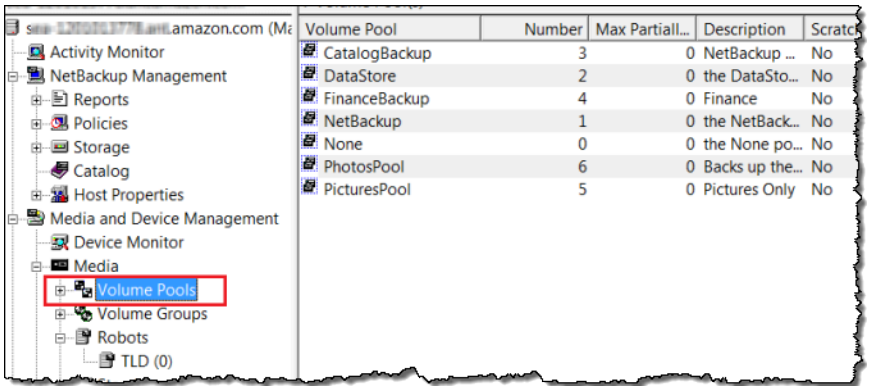


### To add virtual tapes to a volume pool

1. Expand the **Robots** node, and select the **TLD(0)** robot to display the virtual tapes this robot is aware of.

   Note that if you have previously connected a robot, your gateway-VTL robot might have a different name.
2. From the list of virtual tapes, right-click the tape you want to add to the volume pool, and click **Change** to open the **Change Volumes** dialog box. The following screenshot shows the **Change Volumes** dialog box.

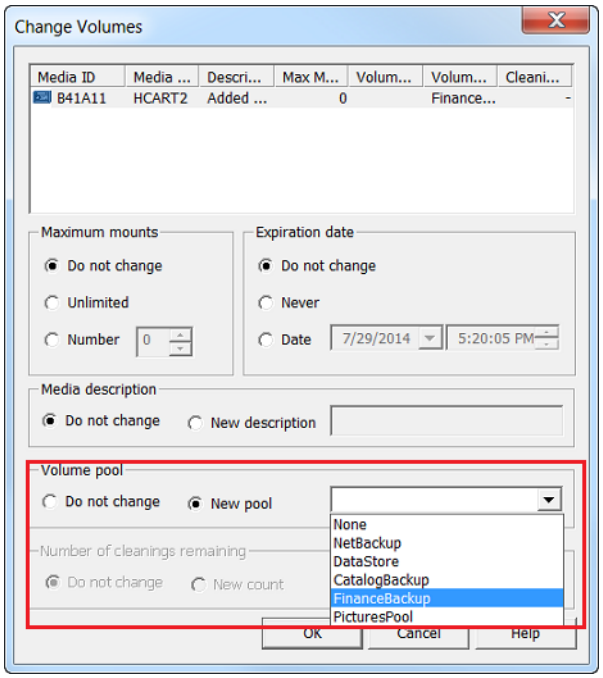


3. In the **Volume Pool** category, select **New pool**.
4. From the **New pool** list, select the pool you just created, and then click **OK**.

   You can verify that your volume pool contains the virtual tape that you just added by expanding the **Media** node and clicking your volume pool.

### To create a backup policy

The backup policy specifies what data to back up, when to back it up, and which volume pool to use.

1. Click your **Master Server** to return to the Symantec NetBackup console.

   The following screenshot shows the NetBackup console with **Create a Policy** selected.

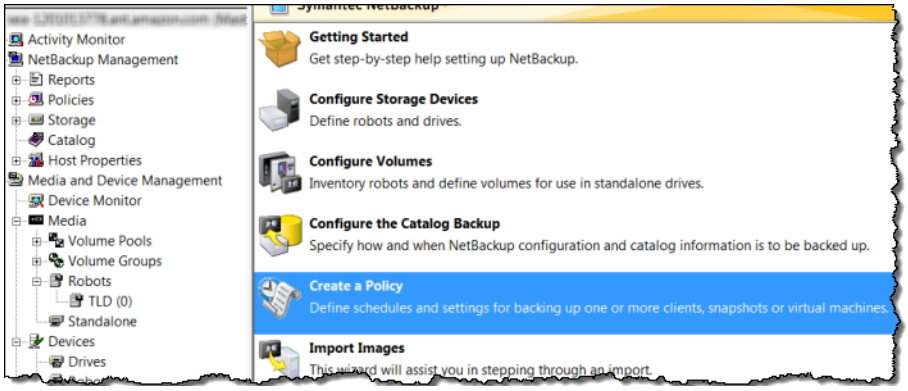


2. Click **Create a Policy** to open the **Policy Configuration Wizard** window.
3. Select **File systems, databases, applications**, and click **Next**.
4. In the **Policy Name** box, type a name for your policy and verify that **MS-Windows** is selected from the **Select the policy type** list, and then click **Next**.
5. In the **Client List** window, click **Add**, type the host name of your computer in the **Name** column, and then click **Next**. This step applies the policy you are defining to localhost (your client computer).

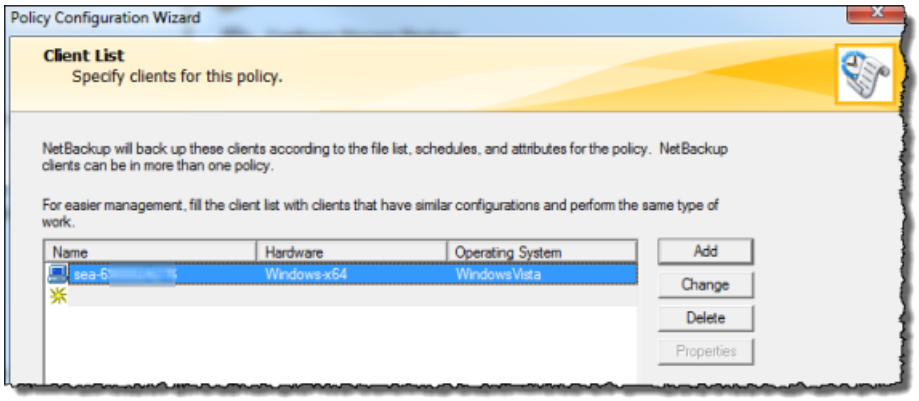
Policy Configuration Wizard
Client List
Specify clients for this policy.
NetBackup will back up these clients according to the file list, schedules, and attributes for the policy. NetBackup clients can be in more than one policy.
For easier management, fill the client list with clients that have similar configurations and perform the same type of work.
Name
Hardware
Operating System
Windows-x64
WindowsVista
Add
Change
Delete
Properties

6. In the **Files** window, click **Add**, and then click the folder icon.

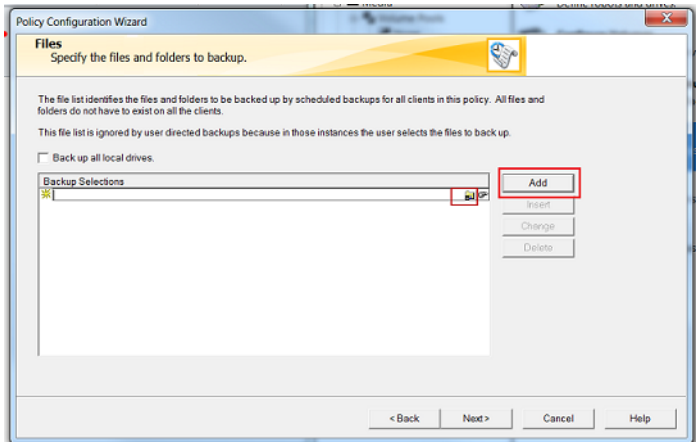
Policy Configuration Wizard
Files
Specify the files and folders to backup.
The file list identifies the files and folders to be backed up by scheduled backups for all clients in this policy. All files and folders do not have to exist on all the clients.
This file list is ignored by user directed backups because in those instances the user selects the files to back up.
Back up all local drives.
Backup Selections
Add
Insert
Change
Delete
< Back
Next >
Cancel
Help

7. In the **Browse** window, browse to the folder or files you want to back up, click **OK**, and then click **Next**.
8. In the **Backup Types** window, accept the defaults, and then click **Next**.

   **Note**
   If you want to initiate the backup yourself, select **User Backup**.

9. In the **Frequency and Retention** window, select the frequency and retention policy you want to apply to the backup. For this exercise, you can accept all the defaults and click **Next**.

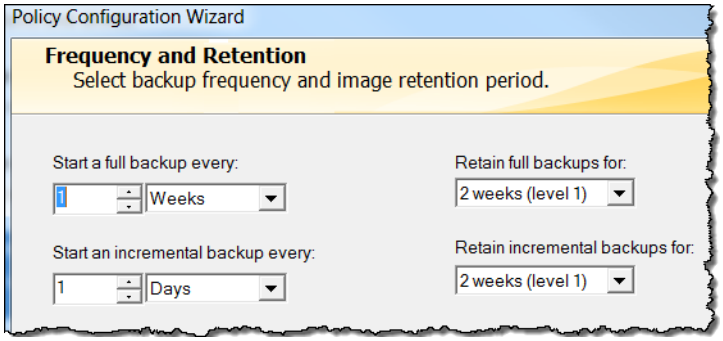


10. In the **Start** window, select **Off hours**, and then click **Next**. This selection specifies that your folder should be backed up during off hours only.

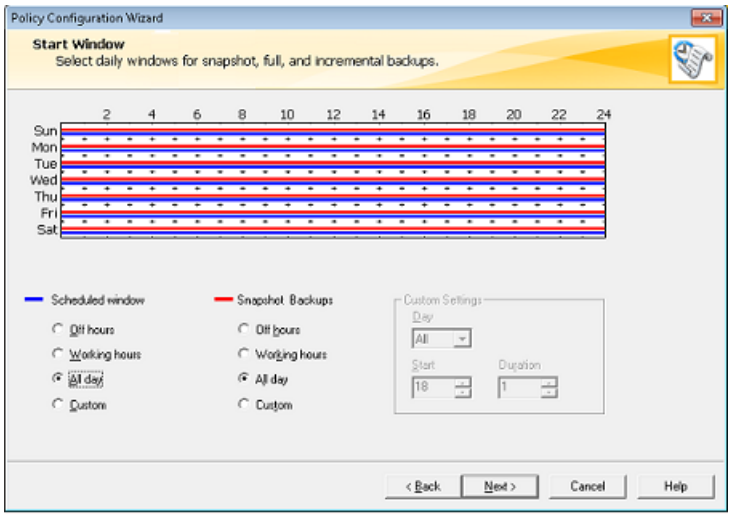


11. In the **Policy Configuration** wizard, click **Finish**.

The policy runs the backups according to the schedule. You can also perform a manual backup at any time, which we will do in the next step.

### To perform a manual backup

1. In the left navigation pane of the NetBackup console, expand the **NetBackup Management** node.
2. Expand the **Policies** node.
3. Right-click your policy, and click **Manual Backup**.

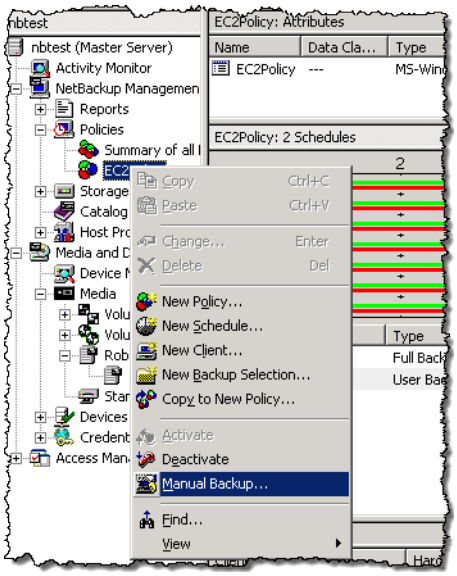


4. In the **Manual Backup** window, select a schedule, select a client, and then click **OK**.

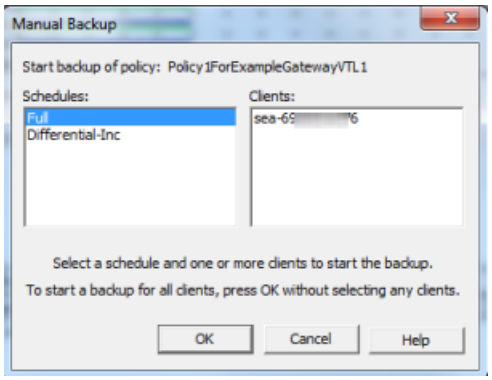


5. In the **Manual Backup Started** dialog box that appears, click **OK**.
6. In the left navigation pane, click **Activity Monitor** to view the status of your backup in the **Job ID** column.

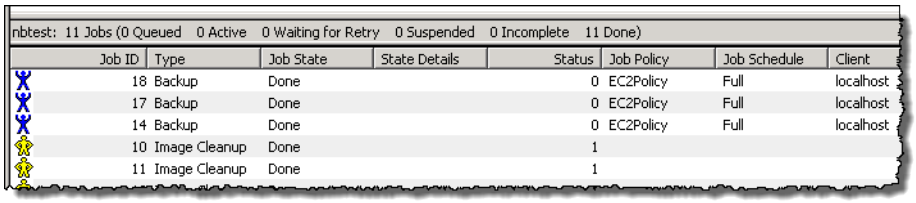
nbtest: 11 Jobs (0 Queued 0 Active 0 Waiting for Retry 0 Suspended 0 Incomplete 11 Done)
Job ID Type Job State State Details Status Job Policy Job Schedule Client
18 Backup Done 0 EC2Policy Full localhost
17 Backup Done 0 EC2Policy Full localhost
14 Backup Done 0 EC2Policy Full localhost
10 Image Cleanup Done 1
11 Image Cleanup Done 1

To find the barcode of the virtual tape where NetBackup wrote the file data during the backup, look in the **Job Details** window as described following. You will need this barcode in the procedure in the next section, where you archive the tape to the virtual tape shelf (VTS).

### To find the barcode of a tape

1. In **Activity Monitor**, right-click the identifier of your backup job in the **Job ID** column, and then click **Details**.
2. In the **Job Details** window, click the **Detailed Status** tab.
3. In the **Status** box, locate the media ID. For example, in the following screenshot, the media ID is **87A222**. This ID helps you determine which tape you have written data to.

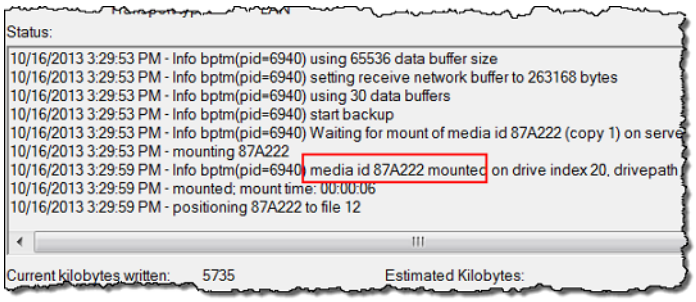
Status:
10/16/2013 3:29:53 PM - Info bptm(pid=6940) using 65536 data buffer size
10/16/2013 3:29:53 PM - Info bptm(pid=6940) setting receive network buffer to 263168 bytes
10/16/2013 3:29:53 PM - Info bptm(pid=6940) using 30 data buffers
10/16/2013 3:29:53 PM - Info bptm(pid=6940) start backup
10/16/2013 3:29:53 PM - Info bptm(pid=6940) Waiting for mount of media id 87A222 (copy 1) on serve
10/16/2013 3:29:53 PM - mounting 87A222
10/16/2013 3:29:59 PM - Info bptm(pid=6940) media id 87A222 mounted on drive index 20, drivepath
10/16/2013 3:29:59 PM - mounted; mount time: 00:00:06
10/16/2013 3:29:59 PM - positioning 87A222 to file 12
Current kilobytes written: 5735
Estimated Kilobytes:

You have now successfully deployed a gateway-VTL, created virtual tapes, and backed up your data. Next, you can archive the virtual tapes and retrieve them from the VTS.

## Archiving the Tape

When you archive a tape, AWS Storage Gateway moves the tape from your gateway's virtual tape library (VTL) to the virtual tape shelf (VTS), which provides offline storage. You initiate tape archival by ejecting the tape using your backup application.

### To archive a virtual tape

1. In the NetBackup Administration console, expand the **Media and Device Management** node, and expand the **Media** node.
2. Expand **Robots** and click **TLD**(0).
3. Right-click the virtual tape you want to archive, and click **Eject Volume From Robot**.

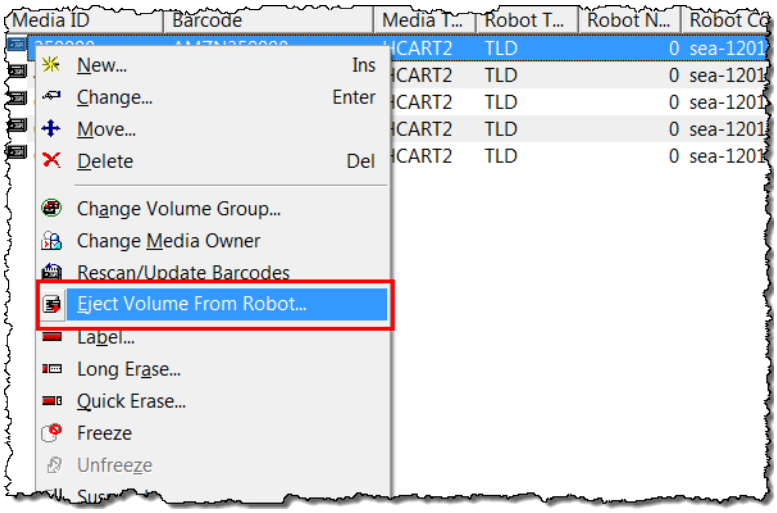


4. In the **Eject Volumes** window, make sure the **Media ID** matches the virtual tape you want to eject, and then click **Eject**.

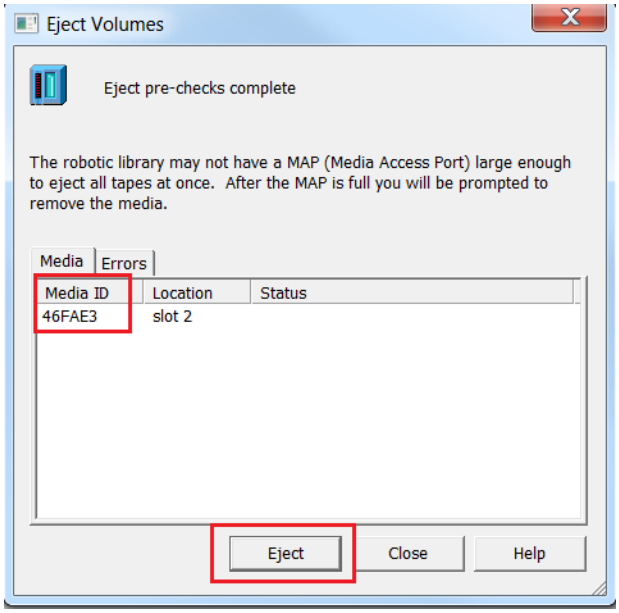


5. In the dialog box, click **Yes**. The dialog box is shown following.

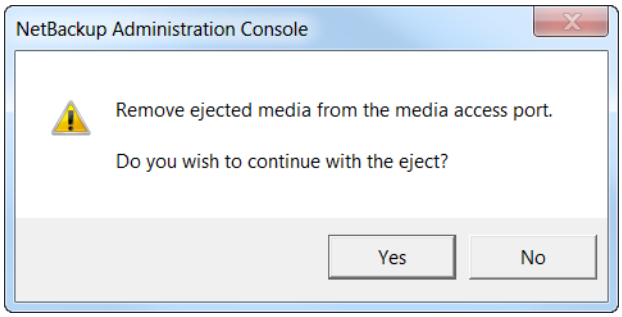
NetBackup Administration Console
Remove ejected media from the media access port.
Do you wish to continue with the eject?
Yes
No

When the eject process is completed, the status of the tape in the **Eject Volumes** dialog box indicates that the eject succeeded.

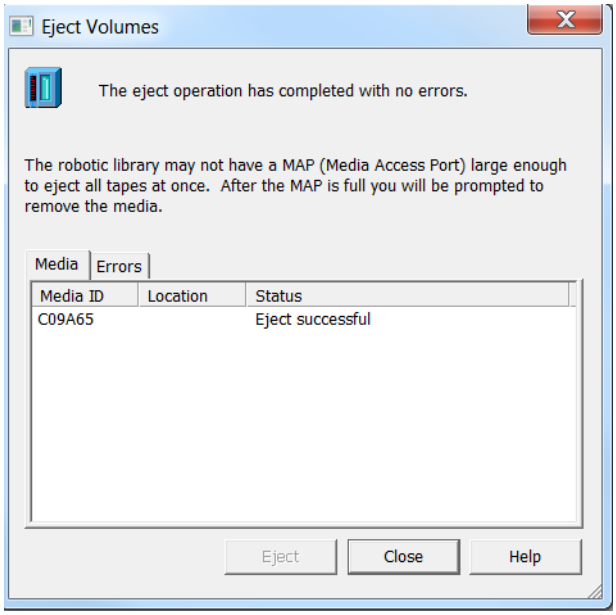
Eject Volumes
The eject operation has completed with no errors.
The robotic library may not have a MAP (Media Access Port) large enough to eject all tapes at once. After the MAP is full you will be prompted to remove the media.
Media
Errors
Media ID
Location
Status
C09A65
Eject successful
Eject
Close
Help

6. Click **Close** to close the **Eject Volumes** window.
7. In the AWS Storage Gateway console, verify the status of the tape you are archiving in the gateway's VTL. It can take some time to finish uploading data to AWS. During this time, the ejected tape will be listed in the gateway's VTL with the status IN TRANSIT TO VTS. When archiving starts, the status will be ARCHIVING. Once data upload has completed, the ejected tape will no longer be listed in the VTL.
8. To verify that the virtual tape is no longer listed in your gateway, click your gateway, and then click **VTL Tape Cartridges**.
9. In the navigation pane of the AWS Storage Gateway console, click the VTS. Verify that your archived tape is listed in the VTS and that its status is ARCHIVED.

## Retrieving the Archived Tape from the VTS Back to Your Gateway-VTL

Archived tapes are stored in a virtual tape shelf (VTS), which provides offline storage. If you want to access tape data, you must first retrieve the tape from the VTS back to your gateway. In this step, you will retrieve the tape that you archived in the preceding step.

**Note**
It takes about 24 hours to retrieve a tape from the VTS to a gateway.

### To retrieve an archived tape

1. In the navigation pane of the AWS Storage Gateway console, click **Virtual Tape Shelf**. The console displays all virtual tapes that have been archived by all your gateways.
2. Select the virtual tape you want to retrieve, and click **Retrieve Tape**.

   **Note**
   The status of the virtual tape you want to retrieve must be ARCHIVED.

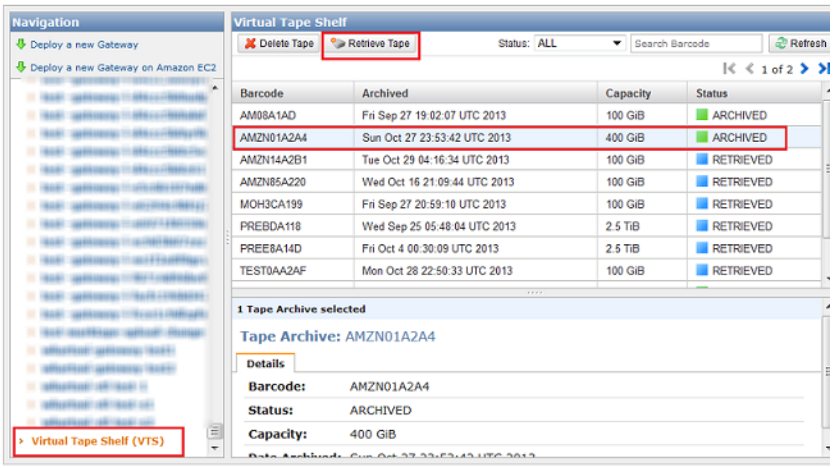


3. In the **Tape Barcode** field of the **Retrieve Tape** wizard, verify that the barcode identifies the virtual tape you want to retrieve.

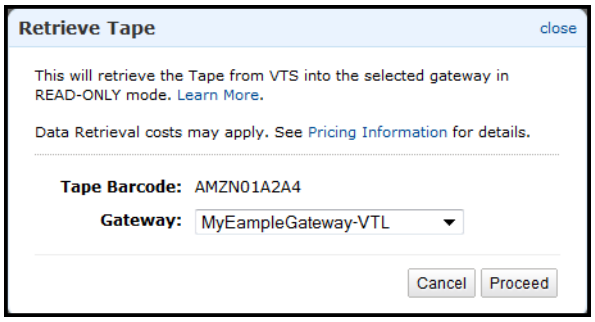
Retrieve Tape
close
This will retrieve the Tape from VTS into the selected gateway in READ-ONLY mode. Learn More.
Data Retrieval costs may apply. See Pricing Information for details.
Tape Barcode: AMZN01A2A4
Gateway: MyEampleGateway-VTL
Cancel
Proceed

4. In the **Gateway** list, select the gateway you want to retrieve the archived tape to, and then click **Proceed**.

The status of the tape changes from ARCHIVED to RETRIEVING. After all the data is moved, the status of the virtual tape in the VTS changes to RETRIEVED, and the tape appears in your gateway's VTL.

**Note**
Retrieved virtual tapes are read-only.

## Restoring Data from the Tape

In this step, you restore data from the virtual tape to your client computer.

**Note**
This step does not use the NetBackup Administration Console that you used in the previous steps. Instead, you will use the Backup, Archive, and Restore software installed with Symantec NetBackup software.

### To restore data

1. Start the Backup, Archive, and Restore software, and run it as an administrator.
2. Click the **Select for Restore** tab.

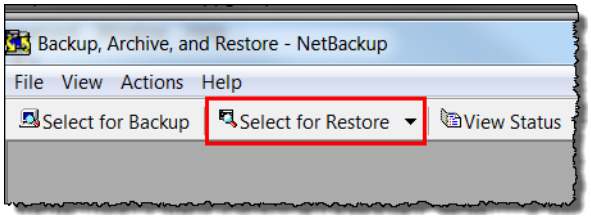


3. If you have previously backed up data, you will see backup icons in the **NetBackup History** pane. In this example, there is only one backup icon.

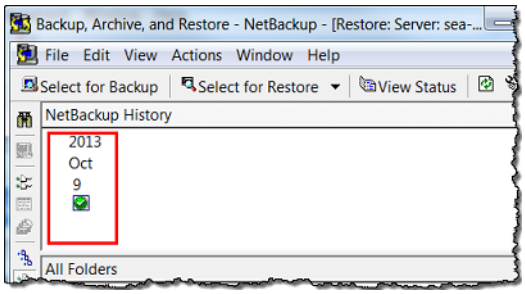


4. Select the backup icon that represents the backup you want to restore.
5. In the **All Folders** pane, select the folder you want to restore, and then click the **Restore Marked Files** icon in the left pane. This icon is not labeled. The name appears when you rest your mouse on the icon.

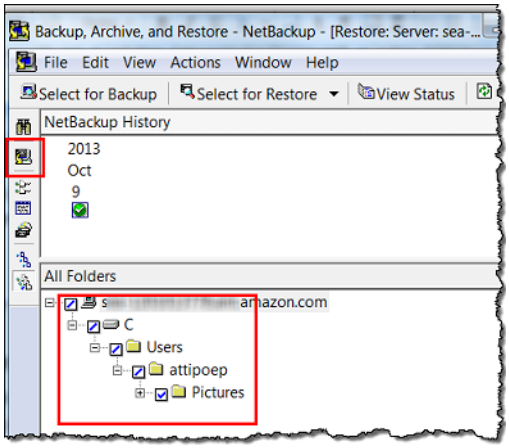


6. In the **Restore Marked Files** window, select the **Restore everything to a different location (maintaining existing structure)** button. This selection avoids overwriting your original data.

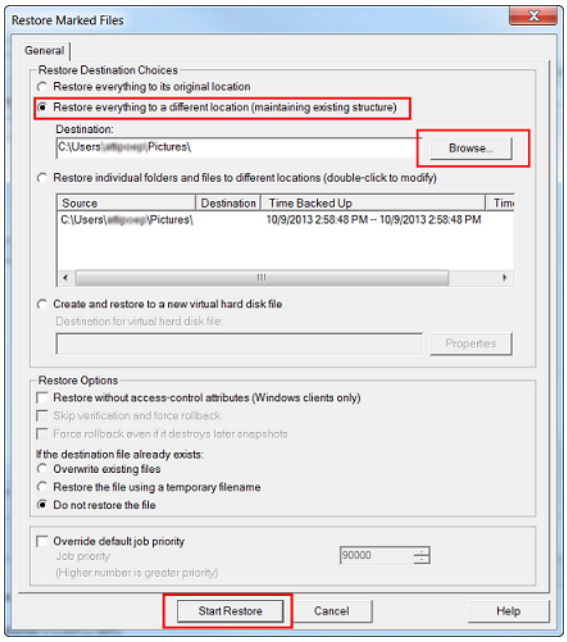
Restore Marked Files
General
Restore Destination Choices
Restore everything to its original location
Restore everything to a different location (maintaining existing structure)
Destination:
Browse...
Restore individual folders and files to different locations (double-click to modify)
Source
Destination
Time Backed Up
10/9/2013 2:58:48 PM -- 10/9/2013 2:58:48 PM
Create and restore to a new virtual hard disk file
Properties
Restore Options
Restore without access-control attributes (Windows clients only)
Skip verification and force rollback
Force rollback even if it destroys later snapshots
If the destination file already exists:
Overwrite existing files
Restore the file using a temporary filename
Do not restore the file
Override default job priority
Job priority
90000
(Higher number is greater priority)
Start Restore
Cancel
Help

7. In the **Destination** box, browse to the folder you want to restore the data to, and then click **Start Restore**.
8. In the dialog box that appears, click **Yes** to view the progress of the restore process.

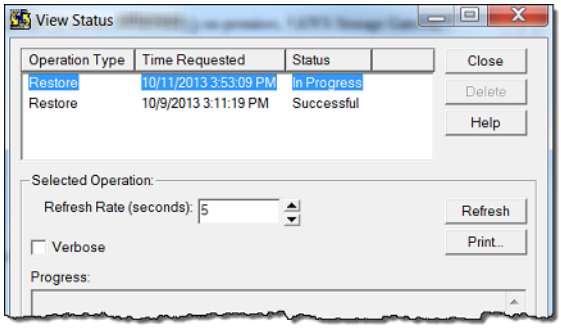


In the **View Status** window, you can see the status of the restore process. If the restore succeeds, the status changes to **Successful**.

**Next Step**

Step 6: Clean Up After the Example Exercise (p. 275)

# Testing Your Setup by Using Symantec Backup Exec

AWS Storage Gateway-VTL supports backing up data to virtual tapes and archiving the tapes in the virtual tape shelf (VTS) by using Symantec Backup Exec. In this topic, we provide basic documentation needed to perform backup using the following versions of Backup Exec:

- Symantec Backup Exec 2014
- Symantec Backup Exec 15

The procedure for using these versions of Backup Exec with AWS Storage Gateway-VTL is the same. For detailed information about how to use Backup Exec, go to the How to Create Secure Backups with Backup Exec video on the Symantec website. For Symantec support information on hardware compatibility,

go to the Software Compatibility Lists (SCL), Hardware Compatibility Lists (HCL), and Administrator Guides for Backup Exec (all versions) on the Symantec website. For information about best practices, go to Best Practices for using Symantec Backup products (NetBackup, Backup Exec) with the Amazon Web Services (AWS) Storage Gateway-VTL on the Symantec website.

> **Note**
> You can use Symantec Backup Exec 2012 with AWS Storage Gateway-VTL. For information, see Appendix I: Testing Your Setup by Using Symantec Backup Exec 2012 (p. 426).

This topic shows you how to configure storage, write data to a tape, archive a tape in the VTS, and restore the data.

**Topics**


- Configuring Storage in Backup Exec (p. 264)
- Importing a Tape in Backup Exec (p. 265)
- Writing Data to a Tape in Backup Exec (p. 266)
- Archiving a Tape Using Backup Exec (p. 266)
- Restoring Data from a Tape Archived in Backup Exec (p. 267)
- Disabling a Tape Drive in Backup Exec (p. 267)


## Configuring Storage in Backup Exec

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Backup Exec storage to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245).

### To configure storage

1. Start the Backup Exec software, and then click the yellow icon in top-left corner on the toolbar.
2. Click **Configuration and Settings**, and then click **Backup Exec Services** to open the Backup Exec Service Manager.

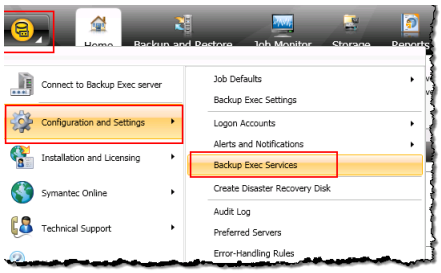


3. Click **Restart All Services**. Backup Exec recognizes the VTL devices (that is, the medium changer and tape drives). The restart process might take a few minutes.

> **Note**
> AWS Storage Gateway-VTL provides 10 tape drives. However, your Backup Exec license agreement might require your backup software to work with fewer than 10 tape drives. In that case, you must disable tape drives in the Backup Exec robotic library to leave only the number of tape drives allowed by your license agreement enabled. For instructions, see Disabling a Tape Drive in Backup Exec (p. 267).

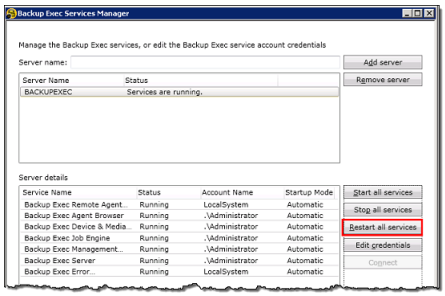


4. After the restart is completed, close the Backup Exec Service Manager.

## Importing a Tape in Backup Exec

You are now ready to import a tape from your gateway into a slot.

1. Click the **Storage** tab, and then expand the **Robotic library** tree to display the VTL devices.

   **Important**
   Symantec Backup Exec software requires the AWS Gateway-VTL medium changer type. If the medium changer type listed under **Robotic library** is not AWS Gateway-VTL, you must change it before you configure storage in the backup software. For information about how to select a different medium changer type, see Selecting a Medium Changer After Gateway Activation (p. 424).

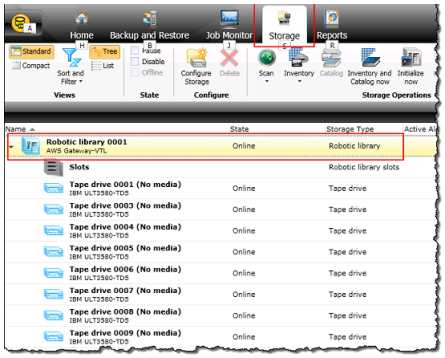


2. Double-click the **Slots** icon to display all slots.

   **Note**
   When you import tapes into the robotic library, the tapes are stored in slots instead of tape drives. Therefore, the tape drives might have a message that indicates there is no media in the drives (No media). When you initiate a backup or restore job, the tapes will be moved into the tape drives.
   You must have tapes available in your gateway tape library to import a tape into a storage slot. For instructions on how to create tapes, see Creating Virtual Tapes in Your Gateway-VTL (p. 279).

3. Right-click an empty slot, select **Import**, and then click **Import media now** in the submenu. In the following screenshot, slot number **3** is empty. You can select more than one slot and import multiple tapes in a single import operation.

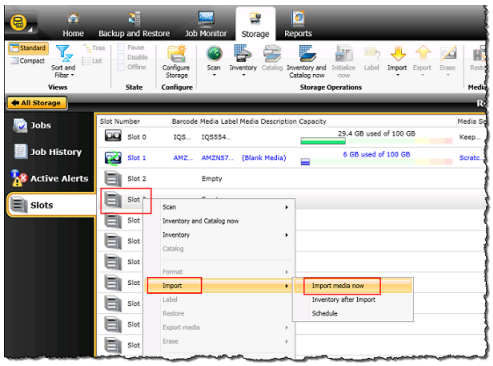


4. In the **Media Request** window that appears, click the **View details** link.

5. In the **Action Alert: Media Intervention** window, click **Respond OK** to insert the media into the slot.

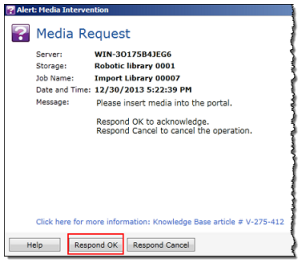
Alert: Media Intervention
Media Request
Server: WIN-3O17SB4JEG6
Storage: Robotic library 0001
Job Name: Import Library 00007
Date and Time: 12/30/2013 5:22:39 PM
Message: Please insert media into the portal.
Respond OK to acknowledge.
Respond Cancel to cancel the operation.
Click here for more information: Knowledge Base article # V-275-412
Help
Respond OK
Respond Cancel

The tape appears in the slot you selected.

**Note**
Tapes that are imported include empty tapes and tapes that have been retrieved from the VTS to the gateway.

## Writing Data to a Tape in Backup Exec

You write data to a gateway-VTL virtual tape by using the same procedure and backup policies you do with physical tapes. For detailed information, go to the *Backup Exec Administrative Guide* in the documentation section in the Backup Exec software.

## Archiving a Tape Using Backup Exec

When you archive a tape, AWS Storage Gateway moves the tape from your gateway's virtual tape library (VTL) to the VTS—that is, the offline storage. You begin tape archival by exporting the tape using your Backup Exec software.

### To archive your tape in VTS

1. Click the **Storage** menu, double-click **Slots**, right-click the slot you want to export the tape from, click **Export media**, and then click **Export media now**. You can select more than one slot and export multiple tapes in a single export operation.

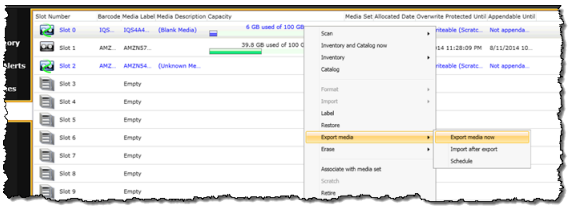

2. In the **Media Request** pop-up window, click **View details**, and then click **Respond OK** in the **Alert: Media Intervention** window.

   In the AWS Storage Gateway console, you can verify the status of the tape you are archiving. It might take some time to finish uploading data to AWS. During this time, the exported tape will be listed in the gateway's VTL with the status IN TRANSIT TO VTS. When the upload is completed and the archiving process begins, the status changes to ARCHIVING. When data archiving has completed, the exported tape will no longer be listed in the VTL.
3. Click your gateway, and then click **VTL Tape Cartridges** and verify that the virtual tape is no longer listed in your gateway.
4. On the Navigation pane of the AWS Storage Gateway console, click **Virtual Tape Shelf (VTS)**. Verify that your archived tape is listed in the VTS and that the status is ARCHIVED.

## Restoring Data from a Tape Archived in Backup Exec

Restoring your archived data is a two-step process.

### To restore data from an archived tape

1. Retrieve the archived tape from your VTS to a gateway-VTL. For instructions, see Retrieving the Archived Tape from the VTS Back to Your Gateway-VTL (p. 260).
2. Use Backup Exec to restore the data. This process is the same as restoring data from physical tapes. For instructions, go to the *Backup Exec Administrative Guide* in the documentation section in the Backup Exec software.

## Disabling a Tape Drive in Backup Exec

AWS Storage Gateway-VTL provides 10 tape drives, but you might decide to use fewer tape drives. In that case, you disable the tape drives you will not use.

1. Open Backup Exec, and click the **Storage** tab.
2. In the **Robotic library** tree, right-click the tape drive you want to disable, and then click **Disable**.

# Testing Your Setup by Using Microsoft System Center 2012 R2 Data Protection Manager

AWS Storage Gateway-VTL supports backing up data to virtual tapes and archiving the tapes in a virtual tape shelf (VTS) by using Microsoft System Center 2012 R2 Data Protection Manager (DPM). In this topic, we provide basic documentation on how to configure the DPM backup software for a gateway–virtual tape library (VTL) and perform a backup. For detailed information about how to use DPM, go to the DPM documentation on the Microsoft System Center website. This topic shows you to configure storage, write data to a tape, archive a tape in the VTS, and restore the data.

### Topics

- Configuring DPM to Recognize VTL Devices (p. 268)
- Importing a Tape into DPM (p. 269)

- Writing Data to a Tape in DPM (p. 270)
- Archiving a Tape by Using DPM (p. 270)
- Restoring Data from a Tape Archived in DPM (p. 270)

## Configuring DPM to Recognize VTL Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure DPM to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245).

By default, the DPM server does not recognize gateway-VTL devices. To configure the server to work with the gateway-VTL devices, you perform the following tasks:

1. Update the device drivers for the VTL devices to expose them to the DPM server.
2. Manually map the VTL devices to the DPM tape library.

### To update the VTL device drivers

- In Device Manager, update the driver for the medium changer. For instructions, see Updating the Device Driver for Your Medium Changer (p. 425).

You use the DPMDriveMappingTool to map your VTL tape drives to the DPM tape library.

### To map VTL tape drives to the DPM server tape library

1. Create at least one tape for your gateway. For information on how to do this on the console, see Step 3: Create Virtual Tapes Using the AWS Storage Gateway Console (p. 243).
2. Import the tape into the DPM library. For information on how to do this, see Importing a Tape into DPM (p. 269).
3. If the DPMLA service is running, stop it by opening a command terminal and typing the following on the command line:

   ```
   net stop DPMLA
   ```
4. Locate the following file on the DPM server: `%ProgramFiles%\System Center 2012 R2\DPM\DPM\Config\DPMLA.xml`.

   **Note**
   If this file exists, the DPMDriveMappingTool will overwrite it. If you want to preserve your original file, create a backup copy.
5. Open a command terminal, change the directory to `%ProgramFiles%\System Center 2012 R2\DPM\DPM\Bin`, and run the following command:

   ```
   C:\Microsoft System Center 2012 R2\DPM\DPM\bin>DPMDriveMappingTool.exe
   ```

   The output for the command looks like the following:

   ```
   Performing Device Inventory ...
   Mapping Drives to Library ...
   ```

```
Adding Standalone Drives ...
Writing the Map File ...
Drive Mapping Completed Successfully.
```

## Importing a Tape into DPM

You are now ready to import tapes from your gateway-VTL into the DPM backup software library.

### To import tapes into the DPM backup software library

1. On the DPM server, open the Management Console, click **Rescan**, and then click **Refresh**. Doing this displays your medium changer and tape drives.

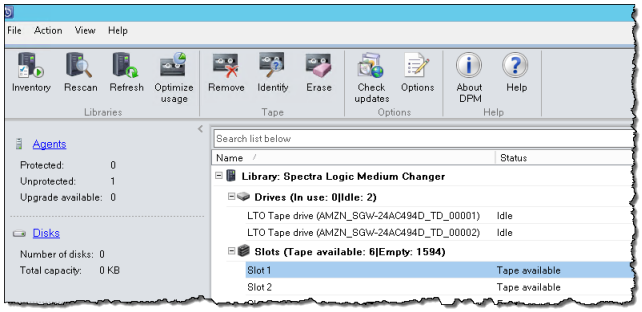


2. Right-click the media changer in the **Library** section, and then click **Add tape (I/E port)** to add a tape to the **Slots** list.

   **Note**
   The process of adding tapes can take several minutes to complete.

   The tape label appears as **Unknown**, and the tape is not usable. For the tape to be usable, you must identify it.

3. Right-click the tape you want to identify, and then click **Identify unknown tape**.

   **Note**
   The process of identifying tapes can take a few seconds or a few minutes.

   When identification is complete, the tape label changes to **Free**. That is, the tape is free for data to be written to it.

   In the following screenshot, the tape in slot 2 has been identified and is free to use but the tape in slot 3 is not.

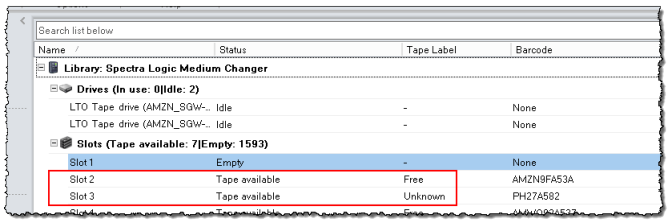
Search list below
Name
Status
Tape Label
Barcode
Library: Spectra Logic Medium Changer
Drives (In use: 0|Idle: 2)
LTO Tape drive (AMZN_SGW-... Idle - None
LTO Tape drive (AMZN_SGW-... Idle - None
Slots (Tape available: 7|Empty: 1593)
Slot 1 Empty - None
Slot 2 Tape available Free AMZN9FA53A
Slot 3 Tape available Unknown PH27A582

## Writing Data to a Tape in DPM

You write data to a gateway-VTL virtual tape by using the same protection procedures and policies you do with physical tapes. You create a protection group and add the data you want to back up, and then back up the data by creating a recovery point. For detailed information about how to use DPM, go to the DPM documentation on the Microsoft System Center website.

## Archiving a Tape by Using DPM

When you archive a tape, AWS Storage Gateway moves the tape from the DPM tape library to the VTS—that is, the offline storage. You begin tape archival by removing the tape from the slot using your backup software—that is, DPM.

### To archive a tape in DPM

1. Right-click the tape you want to archive, and then click **Remove tape (I/E port)**.

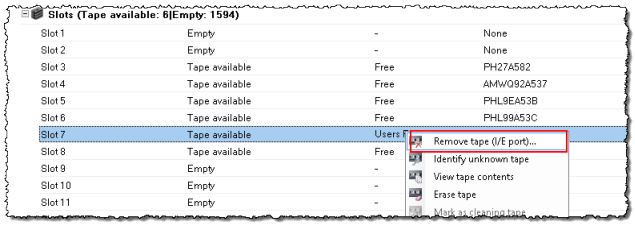
Slots (Tape available: 6|Empty: 1594)
Slot 1 Empty - None
Slot 2 Empty - None
Slot 3 Tape available Free PH27A582
Slot 4 Tape available Free AMWQ92A537
Slot 5 Tape available Free PHL9EA53B
Slot 6 Tape available Free PHL99A53C
Slot 7 Tape available
Slot 8 Tape available Free
Slot 9 Empty -
Slot 10 Empty -
Slot 11 Empty -
Remove tape (I/E port)...
Identify unknown tape
View tape contents
Erase tape

2. In the dialog box that appears, click **Yes**. Doing this ejects the tape from the medium changer's storage slot and moves the tape into one of the gateway's I/E slots. When a tape is moved into the gateway's I/E slot, it is immediately sent to the VTS for archiving.
3. On the AWS Storage Gateway console, click your gateway, and then click **VTL Tape Cartridges** and verify the status of the virtual tape you are archiving.

   The archiving process can take some time to complete. The initial status of the tape is shown as IN TRANSIT TO VTS. When archiving starts, the status changes to ARCHIVING. When archiving is completed, the tape is no longer listed in the VTL.

## Restoring Data from a Tape Archived in DPM

Restoring your archived data is a two-step process.

### To restore data from an archived tape

1. Retrieve the archived tape from your VTS to a gateway-VTL. For instructions, see Retrieving the Archived Tape from the VTS Back to Your Gateway-VTL (p. 260).
2. Use the DPM backup software to restore the data. You do this by creating a recovery point, as you do when restoring data from physical tapes. For instructions, go to Recovering Client Computer Data on the DPM website.

# Testing Your Setup by Using Veeam Backup & Replication

AWS Storage Gateway-VTL supports backing up data to virtual tapes and archiving the tapes in a virtual tape shelf (VTS) by using Veeam Backup & Replication V7 and Veeam Backup & Replication V8. In this topic, we provide basic documentation on how to configure the Veeam Backup & Replication software for a gateway–virtual tape library (VTL) and perform a backup. For detailed information about how to use the Veeam software, go to the Veeam Backup & Replication documentation in the Veeam Help Center. This topic shows you to configure storage, write data to a tape, archive a tape in the VTS, and restore the data.

**Topics**


- Configuring the Veeam Software to Work with VTL Devices (p. 271)
- Importing a Tape into the Veeam Software (p. 272)
- Backing Up Data to a Tape in the Veeam Software (p. 272)
- Archiving a Tape by Using the Veeam Software (p. 272)
- Restoring Data from a Tape Archived in the Veeam Software (p. 273)


## Configuring the Veeam Software to Work with VTL Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Veeam Backup & Replication to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245).

By default, the Veeam software does not recognize gateway-VTL devices. To configure the software to work with gateway-VTL devices, you update the device drivers for the VTL devices to expose them to the Veeam software and then discover the VTL devices.

### Updating VTL Device Drivers

In Device Manager, update the driver for the medium changer. For instructions, see Updating the Device Driver for Your Medium Changer (p. 425).

### Discovering VTL Devices

#### To discover the VTL devices

1. In the Veeam software, click **Backup Infrastructure**. When the gateway-VTL is connected, virtual tapes will be listed in the **Backup Infrastructure** tab.

   **Note**
   Depending on the version of the Veeam Backup & Replication backup software you are using the screenshots in this documentation may look slightly different.

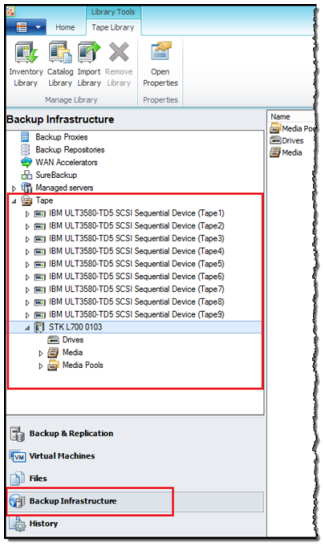


2. Expand the **Tape** tree to see your tape drives and medium changer.
3. Expand the medium changer tree. If your tape drives are mapped to the medium changer, the drives will appear under **Drives**. Otherwise, your tape library and tape drives appear as separate devices.

If the drives are not mapped automatically, follow the instructions on the Veeam website to map the drives.

## Importing a Tape into the Veeam Software

You are now ready to import tapes from your gateway-VTL into the Veeam backup software library.

### To import a tape into the Veeam library

1. Right–click the medium changer, and click **Import** to import the tapes to the I/E slots.
2. Right–click the medium charger, and click **Inventory Library** to identify unrecognized tapes. When you load a new virtual tape into a tape drive for the first time, the tape is not recognized by the Veeam backup application. To identify the unrecognized tape, you inventory the tapes in the tape library.

## Backing Up Data to a Tape in the Veeam Software

Backing data to a tape is a two-step process.

1. You create a media pool and add the tape to the media pool.
2. You write data to the tape.

You create a media pool and write date to a virtual tape by using the same procedures you do with physical tapes. For detailed information about how to back up data, go to the Veeam documentation in the Veeam Help Center.

## Archiving a Tape by Using the Veeam Software

When you archive a tape, AWS Storage Gateway moves the tape from the Veeam tape library to the VTS—that is, the offline storage. You begin tape archival by ejecting from the tape drive to the storage slot and then exporting the tape from the slot to the VTS by using your backup software—that is, the Veeam software.

### To archive a tape in the Veeam library

1. Click **Backup Infrastructure**, and click the media pool that contains the tape you want to archive.

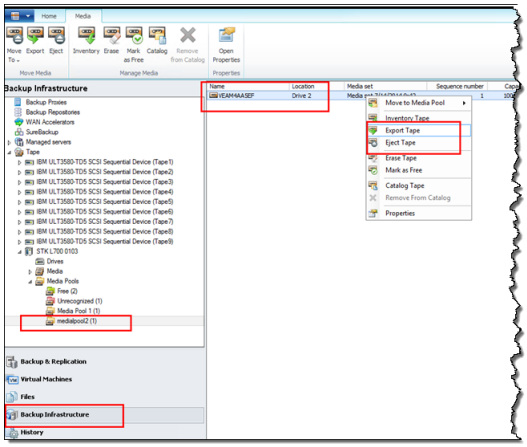


2. Right-click the tape you want to archive, and then click **Eject Tape**.
3. In the **Ejecting tape** box, click **Close**. The location of the tape changes from a tape drive to a slot.
4. Right-click the tape again, and then click **Export**. The status of the tape changes from **Tape drive** to **Offline**.

5. In the **Exporting tape** box, click **Close**. The location of the tape changes from **Slot** to **Offline**.
6. On the AWS Storage Gateway console, click your gateway, and then click **VTL Tape Cartridges** and verify the status of the virtual tape you are archiving.

   The archiving process can take some time to complete. The initial status of the tape is shown as IN TRANSIT TO VTS. When archiving starts, the status changes to ARCHIVING. When archiving is completed, the tape is no longer listed in the VTL.

## Restoring Data from a Tape Archived in the Veeam Software

Restoring your archived data is a two-step process.

### To restore data from an archived tape

1. Retrieve the archived tape from your VTS to a gateway-VTL. For instructions, see Retrieving the Archived Tape from the VTS Back to Your Gateway-VTL (p. 260).
2. Use the Veeam software to restore the data. You do this by creating a restoring a folder file, as you do when restoring data from physical tapes. For instructions, go to Restoring Data from Tape in the Veeam Help Center.

# Testing Your Setup by Using Dell NetVault Backup

AWS Storage Gateway-VTL supports backing up data to virtual tapes and archiving the tapes in a virtual tape shelf (VTS) by using Dell NetVault Backup version 10.0. In this topic, we provide basic documentation on how to configure the NetVault Backup software for a gateway–virtual tape library (VTL) and perform a backup. For detailed information about how to use the NetVault Backup software, go to the NetVault Backup 10.0.1 – Administration Guide. This topic shows you how to configure storage devices, write data to a tape, archive a tape in the VTS, and restore the data. For additional setup information, go to Backing up to Amazon AWS with Dell NetVault Backup on the Dell website.

**Topics**


- Configuring the NetVault Backup Software to Work with VTL Devices (p. 273)
- Backing Up Data to a Tape in the NetVault Backup Software (p. 274)
- Archiving a Tape by Using the NetVault Backup Software (p. 274)
- Restoring Data from a Tape Archived in the NetVault Backup Software (p. 275)


## Configuring the NetVault Backup Software to Work with VTL Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure NetVault Backup to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245).

By default, the NetVault Backup software does not automatically recognize gateway-VTL devices. You must manually add the devices to expose them to the NetVault Backup software and then discover the VTL devices.

### Adding VTL Devices

#### To add the VTL devices

1. In NetVault Backup, click **Mange Devices** in the **Configuration** tab. **Add Devices**.
2. On the Manage Devices page, click **Add Devices**.
3. In the Add Storage Wizard, select **Tape library / media changer** and then click **Next**.

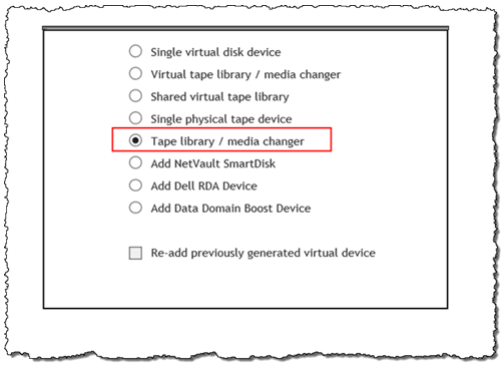
Single virtual disk device
Virtual tape library / media changer
Shared virtual tape library
Single physical tape device
Tape library / media changer
Add NetVault SmartDisk
Add Dell RDA Device
Add Data Domain Boost Device
Re-add previously generated virtual device

4. On the next page, click the client machine that is physically attached to the library and click **Next** to scan for devices.
5. If devices are found they will be displayed. In this case, your medium changer is displayed in the device box.
6. Select your medium changer and click **Next**. Detailed information about the device is displayed in the wizard.
7. On the Add Tapes to Bays page, select **Scan For Devices**, click on your client machine and then click **Next**. All your drives are displayed on the page. NetVault Backup displays the ten bays to which you can add your drives. The bays are displayed one at a time.

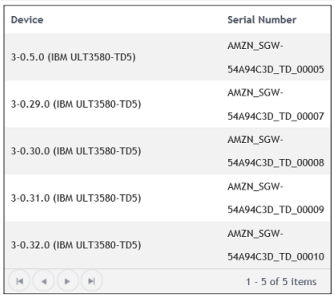

| Device | Serial Number |
|---|---|
| 3-0.5.0 (IBM ULT3580-TD5) | AMZN_SGW-54A94C3D_TD_00005 |
| 3-0.29.0 (IBM ULT3580-TD5) | AMZN_SGW-54A94C3D_TD_00007 |
| 3-0.30.0 (IBM ULT3580-TD5) | AMZN_SGW-54A94C3D_TD_00008 |
| 3-0.31.0 (IBM ULT3580-TD5) | AMZN_SGW-54A94C3D_TD_00009 |
| 3-0.32.0 (IBM ULT3580-TD5) | AMZN_SGW-54A94C3D_TD_00010 |

1 - 5 of 5 Items

8. Click the drive you want to add to the bay that is displayed and click **Next**.

   **Important**
   When you add a drive to a bay, the drive and bay numbers must match. For example, if bay 1 is dispalyed, you must add drive 1. If a drive is not connected, leave its matching bay empty.

9. When your client machine appears, click it and click **Next**. The client machine may appear multiple times.
10. When the drives are displayed, repeat steps 7 through 9 to add all the drives to the bays.
11. In the **Configuration** tab, click **Manage devices** and on the **Manage Devices** page, expand your medium changer to see the devices you added.

## Backing Up Data to a Tape in the NetVault Backup Software

You create a backup job and write data to a virtual tape by using the same procedures you do with physical tapes. For detailed information about how to back up data, go to the Dell NetVault Backup documentation.

## Archiving a Tape by Using the NetVault Backup Software

When you archive a tape, AWS Storage Gateway moves the tape from the NetVault Backup tape library to the VTS— that is, the offline storage. You begin tape archival by ejecting from the tape drive to the storage slot and then exporting the tape from the slot to the VTS by using your backup software—that is, the NetVault Backup software.

### To archive a tape in NetVault Backup

1. In the NetVault Backup Configuration tab, click and expand your medium changer to see your tapes.
2. On the Slots row, click the settings icon to open the **Slots Browser** for the medium changer.

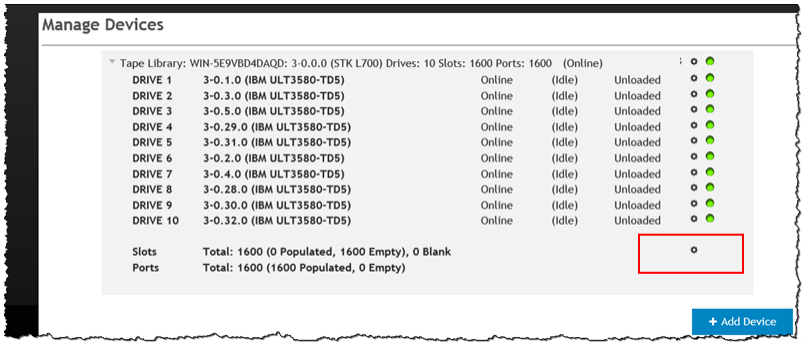
Manage Devices
Tape Library: WIN-5E9VBD4DAQD: 3-0.0.0 (STK L700) Drives: 10 Slots: 1600 Ports: 1600 (Online)
DRIVE 1 3-0.1.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 2 3-0.3.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 3 3-0.5.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 4 3-0.29.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 5 3-0.31.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 6 3-0.2.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 7 3-0.4.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 8 3-0.28.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 9 3-0.30.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
DRIVE 10 3-0.32.0 (IBM ULT3580-TD5) Online (Idle) Unloaded
Slots Total: 1600 (0 Populated, 1600 Empty), 0 Blank
Ports Total: 1600 (1600 Populated, 0 Empty)
+ Add Device

3. In the slots, locate the tape you want to archive, click it, and then click **Export**.

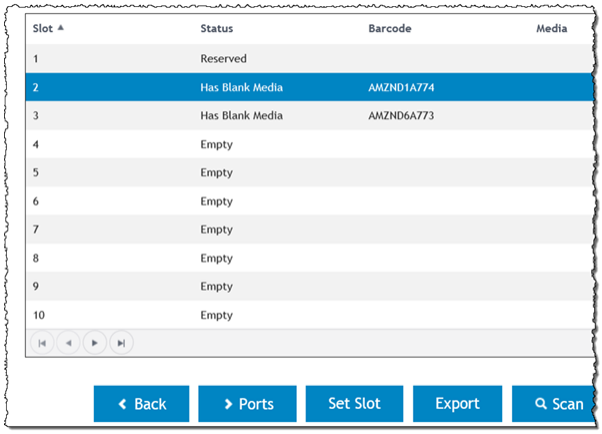


## Restoring Data from a Tape Archived in the NetVault Backup Software

Restoring your archived data is a two-step process.

### To restore data from an archived tape

1. Retrieve the archived tape from your VTS to a gateway-VTL. For instructions, see Retrieving the Archived Tape from the VTS Back to Your Gateway-VTL (p. 260).
2. Use the NetVault Backup software to restore the data. You do this by creating a restoring a folder file, as you do when restoring data from physical tapes. For instructions, go to NetVault Backup 10.0.1 – Administration Guide (Creating a restore job) in the NetVault Backup documentation.

# Step 6: Clean Up After the Example Exercise

After you finish the exercise in this guide, consider cleaning up if you don't plan to continue using your gateway-VTL. Cleaning up lets you avoid incurring charges.

If you plan to continue using your gateway-VTL, see additional information in Where Do I Go from Here? (p. 276)

You should follow these steps in order to ensure you free up the resources successfully.

1. Delete tapes from both your gateway's virtual tape library (VTL) and from the virtual tape shelf (VTS).
   a. Cancel retrieval for any tapes that have the **RETRIEVING** status in your gateway's VTL.

      For instructions, see Canceling Tape Retrieval (p. 281).
   b. Archive any tapes that have the **RETRIEVED** status in your gateway's VTL.

      For instructions, see Archiving Tapes to the VTS (p. 281).
   c. Delete any remaining tapes from your gateway's VTL.

      For instructions, see Deleting Tapes from Your Gateway-VTL (p. 280).
   d. Delete any tapes you have in the VTS.

      For instructions, see Deleting Tapes from the VTS (p. 288).
2. You can keep the gateway-VTL if you plan to use it. Otherwise, delete the gateway.
   a. After deleting all the tapes that have the **RETRIEVING** or **RETRIEVED** status, delete the gateway.

      For instructions, see Deleting Your Gateway-VTL (p. 320).
   b. Delete the AWS Storage Gateway VM from your on-premises host.

## Where Do I Go from Here?

After your gateway-VTL is in production, you can perform several maintenance tasks, such as adding and removing tapes, monitoring and optimizing gateway performance, and troubleshooting. For general information about these management tasks, see Managing Your Activated Gateway-VTL (p. 276).

You can perform some of the gateway-VTL maintenance tasks on the AWS Management Console, such as configuring your gateway's bandwidth rate limits and managing gateway software updates. If your gateway-VTL is deployed on-premises, you can perform some of the maintenance tasks on the gateway's local console, such as routing your gateway-VTL through a proxy and configuring your gateway to use a static IP address. For information about these tasks, see Common Gateway-VTL Maintenance Tasks (p. 302).

If you are running your gateway as an Amazon EC2 instance, you can preform specific maintenance tasks on the Amazon EC2 console, such as adding and removing Amazon EBS volumes. For more information, see Maintenance Tasks When Running a Gateway-VTL on an Amazon EC2 Instance (p. 342).

If you plan to deploy your gateway in production, you should take your real workload into consideration in determining the disk sizes. For information on how to determine real-world disk sizes, see Managing the Upload Buffer and Cache Storage (Gateway-VTL) (p. 304). Also, consider cleaning up if you don't plan to continue using your gateway-VTL. Cleaning up lets you avoid incurring charges. For information on cleanup, see Step 6: Clean Up After the Example Exercise (p. 275).

# Managing Your Activated Gateway-VTL

This section discusses how you can manage your gateway-VTL after you have deployed and activated it. Managing your gateway includes tasks such as configuring cache storage and upload buffer space, working with virtual tapes, doing general maintenance, troubleshooting, and monitoring your gateway's performance. If you have not set up a gateway, see Setting Up Your Gateway-VTL (p. 240).

**Topics**

- Managing Virtual Tapes (p. 277)
- Connecting to the VTL Devices on Your Gateway (p. 288)
- Performing Gateway-VTL Maintenance Tasks (p. 302)

- Monitoring Your Gateway-VTL (p. 351)
- Logging Gateway-VTL API Calls by Using AWS CloudTrail (p. 361)
- Troubleshooting Your Gateway-VTL (p. 363)

# Managing Virtual Tapes

After you deploy a gateway-VTL, you can create tapes. These tapes appear in the virtual tape library (VTL). For each gateway-VTL, there is one VTL. Your backup applications can access tapes in the VTL to read and write data. These tapes are backed by Amazon S3. That is, when you write to these tapes, your gateway stores the data in Amazon S3.

AWS Storage Gateway also provides a virtual tape shelf (VTS) where you can archive tapes in your VTL. A VTS is analogous to an offsite archive location in a real-world tape library.

When you archive a tape, AWS Storage Gateway moves the tape from the VTL to the VTS. AWS Storage Gateway provides one shelf per AWS region for your AWS account. If you set up multiple gateways, they all share the same shelf for tape archival. You can retrieve a tape from the shelf to any gateway-VTL in that region.

> **Note**
> You have one tape library for each gateway-VTL you create in your AWS account, but there is only one virtual shelf per AWS region for your account.

Tapes in VTS are backed by Amazon Glacier. That is, when you archive a tape, AWS Storage Gateway moves tape data from Amazon S3 to Amazon Glacier for long-term archival.

The following section describes various management operations you can perform on tapes that are in the VTL or VTS.

**Topics**

- Managing Tapes in the VTL (p. 277)
- Managing Tapes in the VTS (p. 285)

# Managing Tapes in the VTL

AWS Storage Gateway provides one virtual tape library (VTL) for each gateway-VTL you activate. Initially, the library contains no tapes, but you can create tapes whenever you need to. Your application can read and write to any tapes available on your gateway-VTL. A tape's status must be AVAILABLE for you to write to the tape. These tapes are backed by Amazon S3—that is, when you write to these tapes, the gateway-VTL stores data in Amazon S3. For more information, see Tape Status Information (p. 278).

**Topics**

- Tape Status Information (p. 278)
- Creating Virtual Tapes in Your Gateway-VTL (p. 279)
- Deleting Tapes from Your Gateway-VTL (p. 280)
- Archiving Tapes to the VTS (p. 281)
- Canceling Tape Retrieval (p. 281)
- Canceling Tape Archival (p. 282)
- Recovering a Virtual Tape (p. 282)

The console shows the tapes available for your gateway-VTL in the **VTL Tape Cartridges** tab.

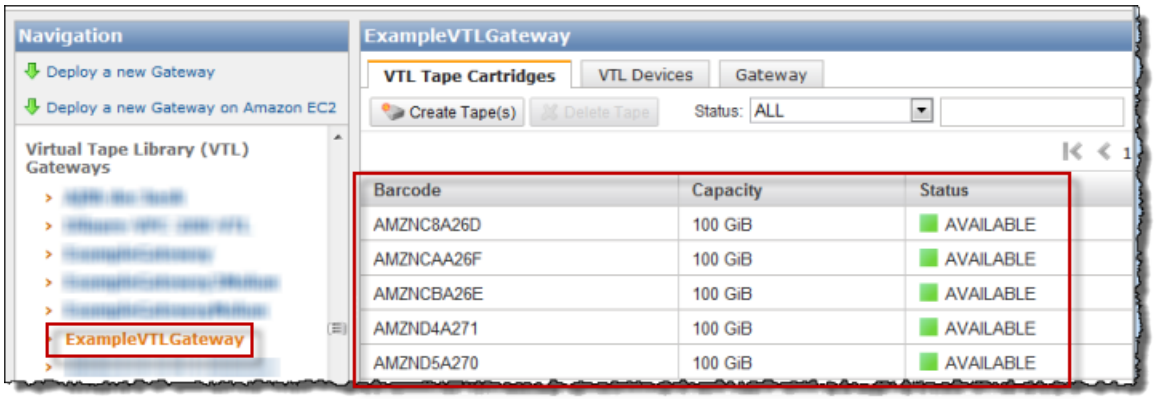


When you have a large number of tapes in the library, the console supports searching for tapes by barcode, by status, or by both. When you search by barcode, you can filter by status.

### To search by barcode and status

1. In the AWS Storage Gateway console, select a gateway-VTL and type the barcode of the virtual tape you want to find in the **Search Barcode** box. By default, AWS Storage Gateway searches for all virtual tapes. However, you can also filter your search by status.
2. In the **Status** drop-down list, select the status of the virtual tapes you want to find. The tapes that match your criteria appear in the AWS Storage Gateway console.

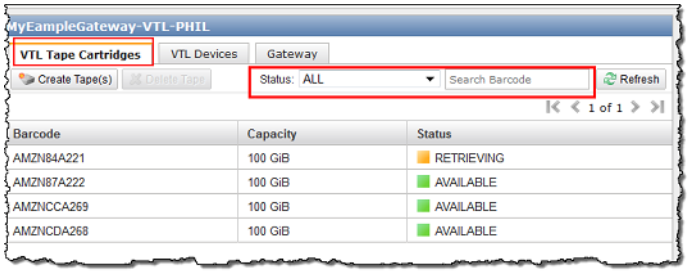


### To search by status

- In the AWS Storage Gateway console, select the status of the tapes you want to find from the **Status** drop-down list. All tapes that have the status that you specify will appear in the AWS Storage Gateway console.

  **Note**
  By default, AWS Storage Gateway displays all tapes regardless of status.

## Tape Status Information

A tape's status must be AVAILABLE for you to read or write to the tape. As shown in the preceding console screenshot, each tape in a VTS has an associated status. The following table lists and describes possible status values.

| Virtual Tape Status | Description | Tape Data Is Stored In |
|---|---|---|
| CREATING | The virtual tape is being created. The tape cannot be loaded into a tape drive, because it is being created. | — |

| Virtual Tape Status | Description | Tape Data Is Stored In |
|---|---|---|
| AVAILABLE | The virtual tape is created and ready to be loaded into a tape drive. | Amazon S3 |
| IN TRANSIT TO VTS | The virtual tape has been ejected and is being uploaded for archive. At this point, your gateway-VTL is uploading data to AWS. If the amount of data being uploaded is small, this status might not be seen. When the upload is completed, the status changes to ARCHIVING. | Amazon S3 |
| ARCHIVING | The virtual tape is being moved by AWS Storage Gateway to the VTS, which is backed by Amazon Glacier. This process happens after the data upload to AWS is completed. | Data is being moved from Amazon S3 to Amazon Glacier |
| DELETING | The virtual tape is being deleted. | Data is being deleted from Amazon S3 |
| DELETED | The virtual tape has been successfully deleted. | — |
| RETRIEVING | The virtual tape is being retrieved from the VTS to your gateway-VTL.<br>**Note**<br>The virtual tape can be retrieved only to a gateway-VTL. | Data is being moved from Amazon Glacier to Amazon S3 |
| RETRIEVED | The virtual tape is retrieved from the VTS. The retrieved tape is write-protected. | Amazon S3 |
| RECOVERED | The virtual tape is recovered and is read-only.<br>When your gateway-VTL is not accessible for any reason, you can recover virtual tapes associated with that gateway-VTL to another gateway-VTL. You must disable the inaccessible gateway-VTL before you can recover the virtual tapes. | Amazon S3 |
| IRRECOVERABLE | The virtual tape cannot be read from or written to. This status indicates an error in your gateway-VTL. | Amazon S3 |

## Creating Virtual Tapes in Your Gateway-VTL

You can create new tapes in your gateway-VTL when you need them. For information about gateway-VTL tape limits, see Appendix F: AWS Storage Gateway Limits (p. 421).

### To create virtual tapes

1. In the AWS Storage Gateway console, select the gateway-VTL for your new tapes, and then click the **VTL Tape Cartridges** tab.
2. Click **Create Tape(s)** to open the **Create Virtual Tapes for Your Storage Gateway** window.

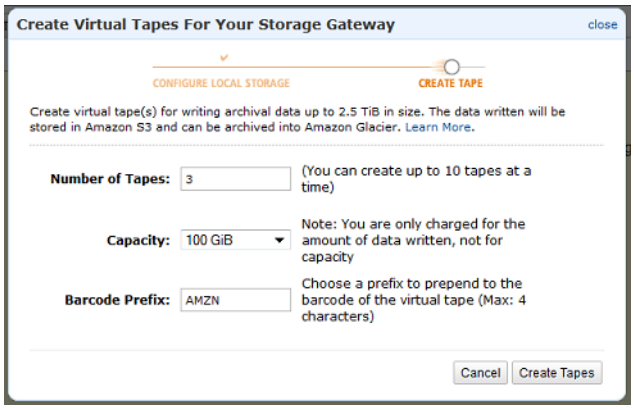
Create Virtual Tapes For Your Storage Gateway
close
CONFIGURE LOCAL STORAGE
CREATE TAPE
Create virtual tape(s) for writing archival data up to 2.5 TiB in size. The data written will be stored in Amazon S3 and can be archived into Amazon Glacier. Learn More.
Number of Tapes:
3
(You can create up to 10 tapes at a time)
Capacity:
100 GiB
Note: You are only charged for the amount of data written, not for capacity
Barcode Prefix:
AMZN
Choose a prefix to prepend to the barcode of the virtual tape (Max: 4 characters)
Cancel
Create Tapes

3. Fill out the form and click **Create Tapes**.

   After the tape status changes to AVAILABLE, your backup application can access the tapes.

## Deleting Tapes from Your Gateway-VTL

You can delete virtual tapes from your gateway-VTL by using the AWS Storage Gateway console.

**Note**
If the tape you want to delete from your gateway-VTL has a status of RETRIEVING, you must first cancel the retrieval before deleting the tape. For more information, see Canceling Tape Retrieval (p. 281). After the tape retrieval is canceled, the tape's status in the VTS changes back to ARCHIVED. You can then delete the tape.
If the tape you want to delete from your gateway-VTL has a status of RETRIEVED, you must first eject the tape using your backup application before deleting the tape. For instructions on how to eject a tape using the Symantec NetBackup software, see Archiving the Tape (p. 259). After the tape is ejected, the tape status in the VTS changes back to ARCHIVED. You can then delete the tape.

### To delete a virtual tape

1. In the AWS Storage Gateway console, select your gateway-VTL, and then click the **VTL Tape Cartridges** tab.
2. Select the virtual tape you want to delete, and click the **Delete Tape** button.

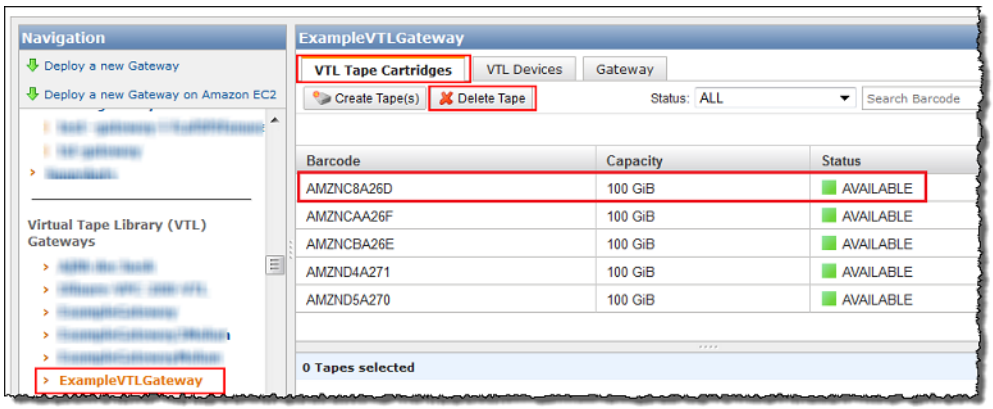


After the tape is deleted, it disappears from the gateway-VTL.

## Archiving Tapes to the VTS

You can archive the virtual tapes that are in your gateway-VTL. When you archive a tape, AWS Storage Gateway moves the tape to the VTS. A VTS is analogous to a real-world off-site tape storage facility. AWS Storage Gateway provides one VTS for each account in an AWS region. If you set up multiple gateways in your account, all those gateways share the same VTS for archiving.

To archive a tape, you use your backup software. Tape archival process consists of three stages, seen as the tape statuses IN TRANSIT TO VTS, ARCHIVING, and ARCHIVED.

- To start the archival process, you eject the tape by using your backup application. After you eject the tape, the tape status becomes IN TRANSIT TO VTS and the tape is no longer accessible to your backup application. In this stage, your gateway-VTL is uploading data to AWS. If needed, you can cancel the archival in progress. For more information about canceling archival, see Canceling Tape Archival (p. 282).
- After the data upload to AWS completes, the tape status changes to ARCHIVING and AWS Storage Gateway begins moving the tape to the VTS. You cannot cancel the archival process at this point.
- After the tape is moved to the VTS, its status changes to ARCHIVED. The tape is now visible in the VTS, and you can retrieve the tape to any of your gateways in the region as the VTS . For more information about tape retrieval, see Retrieving Tapes from the VTS to Your Gateway (p. 286).

The steps involved in archiving a tape depend on your backup software. For instructions on how to archive a tape by using the NetBackup software, see Archiving the Tape (p. 259).

## Canceling Tape Retrieval

After you start the tape retrieval from the VTS to your gateway-VTL, and before the retrieval is complete, you might decide that you don't need to retrieve the tape. You can cancel the retrieval that is in progress. Canceling the retrieval puts the tape back into the VTS.

The following procedure shows you how to cancel the retrieval process.

### To cancel tape retrieval

1. From the AWS Storage Gateway console, click the **VTL Tape Cartridges** tab, and then select the tape you want to stop retrieving.

   For information about how to find a tape in the console, see Managing Tapes in the VTL (p. 277).
2. In the **Details** tab, click **Stop** in the **Status** row.

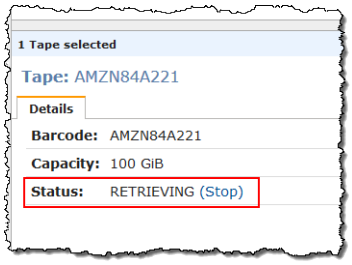


3. In the dialog box that appears, click **OK**. After the retrieval is canceled, AWS Storage Gateway removes the tape from the list of tapes in your VTL, and changes the tape status in VTS from RETRIEVING to ARCHIVED.

   **Note**
   You are charged for the portion the data that was retrieved from VTS before you canceled the tape retrieval.

## Canceling Tape Archival

After you start archiving a tape to the VTS, you might decide you need your tape back. For example, you might want to cancel the archival process, get the tape back because the archival process is taking too long, or read data from the tape. A tape that is being archived goes through three statuses.

- IN TRANSIT TO VTS: Your gateway-VTL is uploading data to AWS.
- ARCHIVING: Data upload is complete and the gateway-VTL is moving the tape to the VTS.
- ARCHIVED: The tape is moved and the VTS and is available for retrieval.

You can cancel archival only when the tape's status is IN TRANSIT TO VTS. Depending on factors such as upload bandwidth and the amount of data being uploaded, this status might or might not be visible in the AWS Storage Gateway console.

### To cancel archival

1. In the AWS Storage Gateway console, click the **VTL Tape Cartridges** tab, and then select the tape you want to stop from archiving. The status of the tape must be IN TRANSIT TO VTS.
2. In the **Details** tab, click **Stop** in the **Status** row.

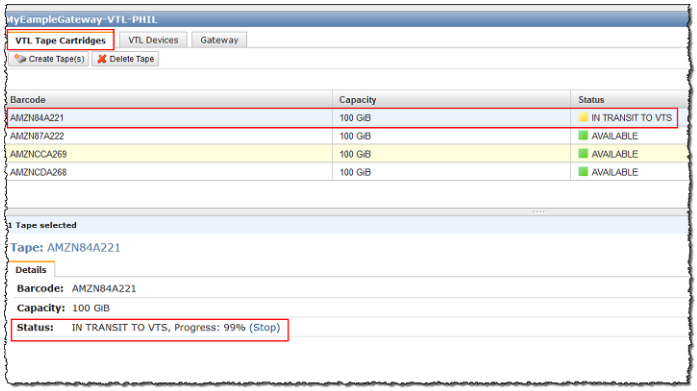


3. Click the **Cancel Archival** button. After archival is canceled, the status of the tape returns to its original status.

## Recovering a Virtual Tape

Although it is rare, your gateway-VTL might encounter an unrecoverable failure. This failure can occur in your hypervisor host, the gateway itself, or the cache disks. If such a failure happens, you can recover your tapes by following the instructions in this section.

### Topics

- Recovering a Virtual Tape from a Malfunctioning Gateway-VTL (p. 282)
- Recovering a Virtual Tape from a Malfunctioning Cache Disk (p. 283)

### Recovering a Virtual Tape from a Malfunctioning Gateway-VTL

If your gateway-VTL or the hypervisor host encounters an unrecoverable failure, you can recover the tapes from that gateway-VTL to another gateway-VTL.

AWS Storage Gateway periodically takes point-in-time snapshots of all the tapes in the library. These snapshots are called recovery points. You can recover tapes to another gateway-VTL from the latest

recovery point. To recover tapes from a failed gateway-VTL to another gateway-VTL, use the following procedure.

### To recover a tape to another gateway-VTL

1. Identify an existing functioning gateway-VTL to serve as your recovery target gateway. If you don't have a gateway-VTL to recover your tapes to, deploy and activate a new gateway-VTL. For information on how to deploy and activate a gateway, see Step 2: Deploy a VM and Activate the Gateway (p. 14).
2. In the AWS Storage Gateway console, click the gateway-VTL you want to disable. All virtual tapes associated with the gateway-VTL are displayed.
3. Disable the failed gateway. This process permanently halts normal function of your gateway-VTL and exposes any available recovery points. For instructions, see Disabling Your Gateway-VTL (p. 319).
4. From the tapes that the disabled gateway displays, select the virtual tape and the recovery point you want to recover. A virtual tape can have multiple recovery points.

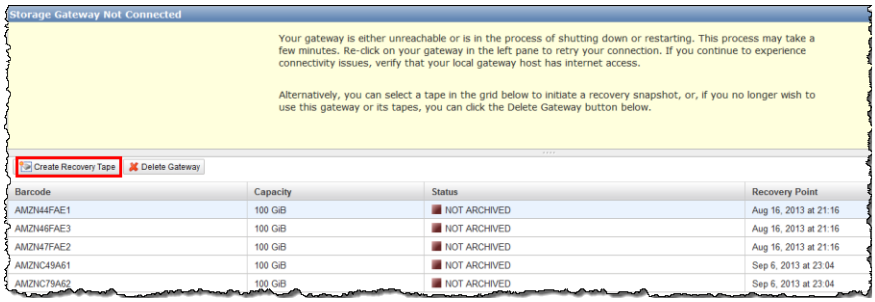
Storage Gateway Not Connected
Your gateway is either unreachable or is in the process of shutting down or restarting. This process may take a few minutes. Re-click on your gateway in the left pane to retry your connection. If you continue to experience connectivity issues, verify that your local gateway host has internet access.
Alternatively, you can select a tape in the grid below to initiate a recovery snapshot, or, if you no longer wish to use this gateway or its tapes, you can click the Delete Gateway button below.
Create Recovery Tape
Delete Gateway
Barcode
Capacity
Status
Recovery Point
AMZN44FAE1
100 GiB
NOT ARCHIVED
Aug 16, 2013 at 21:16
AMZN46FAE3
100 GiB
NOT ARCHIVED
Aug 16, 2013 at 21:16
AMZN47FAE2
100 GiB
NOT ARCHIVED
Aug 16, 2013 at 21:16
AMZNC49A61
100 GiB
NOT ARCHIVED
Sep 6, 2013 at 23:04
AMZNC79A62
100 GiB
NOT ARCHIVED
Sep 6, 2013 at 23:04

5. To begin recovering any tapes you need to the target gateway-VTL, click **Create Recovery Tape**.
6. In the **Create Recovery Tape** dialog box, verify the barcode of the virtual tape you want to recover.

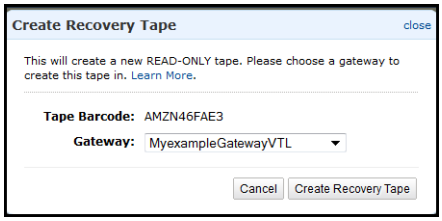
Create Recovery Tape
close
This will create a new READ-ONLY tape. Please choose a gateway to create this tape in. Learn More.
Tape Barcode: AMZN46FAE3
Gateway: MyexampleGatewayVTL
Cancel
Create Recovery Tape

7. In the **Gateway** list, select the gateway-VTL you want to recover the virtual tape to.
8. Click **Create Recovery Tape**.
9. Delete the failed gateway-VTL so you don't get charged. For instructions, see Deleting Your Gateway-VTL (p. 320).

AWS Storage Gateway moves the tape from the failed gateway-VTL to the gateway-VTL you specified. The gateway-VTL marks the tape status as RECOVERED.

## Recovering a Virtual Tape from a Malfunctioning Cache Disk

If your cache disk encounters a error, the gateway prevents read and write operations on virtual tapes in the gateway. For example, an error can occur when a disk is corrupted or removed from the gateway. The AWS Storage Gateway console displays a message about the error and prompts you to take one of two actions that can recover your tapes.

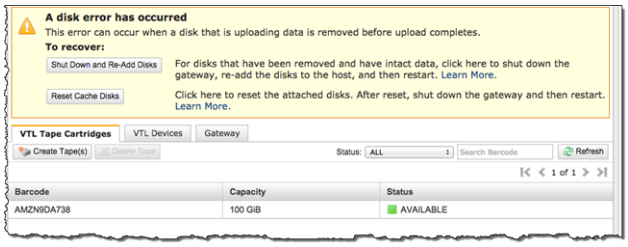
A disk error has occurred
This error can occur when a disk that is uploading data is removed before upload completes.
To recover:
Shut Down and Re-Add Disks
For disks that have been removed and have intact data, click here to shut down the gateway, re-add the disks to the host, and then restart. Learn More.
Reset Cache Disks
Click here to reset the attached disks. After reset, shut down the gateway and then restart. Learn More.
VTL Tape Cartridges
VTL Devices
Gateway
Create Tape(s)
Delete Tape
Status:
ALL
Search Barcode
Refresh
1 of 1
Barcode
Capacity
Status
AMZN9DA738
100 GiB
AVAILABLE

- **Shut Down and Re-Add Disk** – Take this approach if the disk has intact data and has been removed. For example, if the error occurred because a disk was removed from your host by accident but the disk and the data is intact, you can re-add the disk.
- **Reset Cache Disk** – Take this approach if the cache disk is corrupted or not accessible. If the disk error causes the cache disk to be inaccessible, unusable, or corrupted, you can reset the disk. If you reset the cache disk, tapes that have clean data (that is, tapes for which data in the cache disk and Amazon S3 are synchronized) will continue to be available for you to use. However, tapes that have data that is not synchronized with Amazon S3 are automatically recovered. The status of these tapes is set to RECOVERED, but the tapes will be read-only. For information about how to remove a disk from your host, see Adding and Removing Upload Buffer Capacity (Gateway-VTL) (p. 306).

  **Important**
  If the cache disk you are resetting contains data that has not been uploaded to Amazon S3 yet, that data can be lost. After you reset cache disks, there will be no configured cache disks left in the gateway, so you must configure at least one new cache disk for your gateway to function properly.

### To shut down and re-add a disk

1. Shut down the gateway. For information about how to shut down a gateway, see Shutting Down and Starting Your Gateway-VTL (p. 311).
2. Add the disk back to your host, and make sure the disk node number of the disk has not changed. For information about how to add a disk, see Adding and Removing Upload Buffer Capacity (Gateway-VTL) (p. 306).
3. Restart the gateway. For information about how to restart a gateway, see Shutting Down and Starting Your Gateway-VTL (p. 311).

After the gateway restarts, you can verify the status of the cache disks. The status of a disk can be one of the following:

- **present** – The disk is available to use.
- **missing** – The disk is no longer connected to the gateway.
- **mismatch** – The disk node is occupied by a disk that has incorrect metadata, or the disk content is corrupted.

### To reset and reconfigure a cache disk

1. In the **A disk error has occurred** banner, click **Reset Cache Disk**.
2. On the **Configure Your Activated Gateway** page, configure the disk for cache storage. For information about how to do so, see Step 2: Configure Local Storage and Alarms (Gateway-VTL) (p. 241).
3. After you have configured the cache storage, shut down and restart the gateway. The gateway should recover after the restart. You can then verify the status of the cache disk.

### To verify the status of a cache disk

1. On the AWS Storage Gateway console, click your gateway, and then click the **Gateway** tab.
2. Click **Configure Local Storage** to display the **Configure Your Activated Gateway** page. This page shows all local disks in the gateway.

The following screenshot shows an example of the **Configure Your Activated Gateway** page. In this example, the status of the cache disk node xen-vbd-2288 is missing.

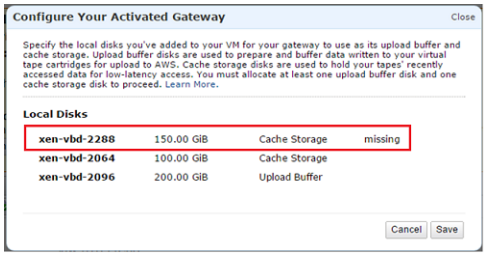
Configure Your Activated Gateway
Close
Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage. Upload buffer disks are used to prepare and buffer data written to your virtual tape cartridges for upload to AWS. Cache storage disks are used to hold your tapes' recently accessed data for low-latency access. You must allocate at least one upload buffer disk and one cache storage disk to proceed. Learn More.
Local Disks
xen-vbd-2288 150.00 GiB Cache Storage missing
xen-vbd-2064 100.00 GiB Cache Storage
xen-vbd-2096 200.00 GiB Upload Buffer
Cancel
Save

**Note**
If you do not complete the recovery process, the gateway displays a banner that prompts you to configure local storage.

## Managing Tapes in the VTS

A virtual tape shelf (VTS) is analogous to an offsite archive location in a real-world tape library. AWS Storage Gateway provides one shelf per AWS region for your AWS account. If you set up multiple gateway-VTLs, they all share the same shelf, and all virtual tapes archived from these gateways go to the same shelf. This functionality allows you to retrieve a virtual tape from the shelf to any gateway-VTL in that region.

**Note**
You have one tape library for each gateway-VTL you create in your AWS account, but there is only one virtual shelf per AWS region in your account.

Tapes in VTS are backed by Amazon Glacier. That is, when you archive a virtual tape to the VTS, AWS Storage Gateway moves tape data from Amazon S3 to Amazon Glacier for long-term archival. You must first retrieve an archived tape before you can access data on the tape, and you can only read the data from an archived tape that has been retrieved.

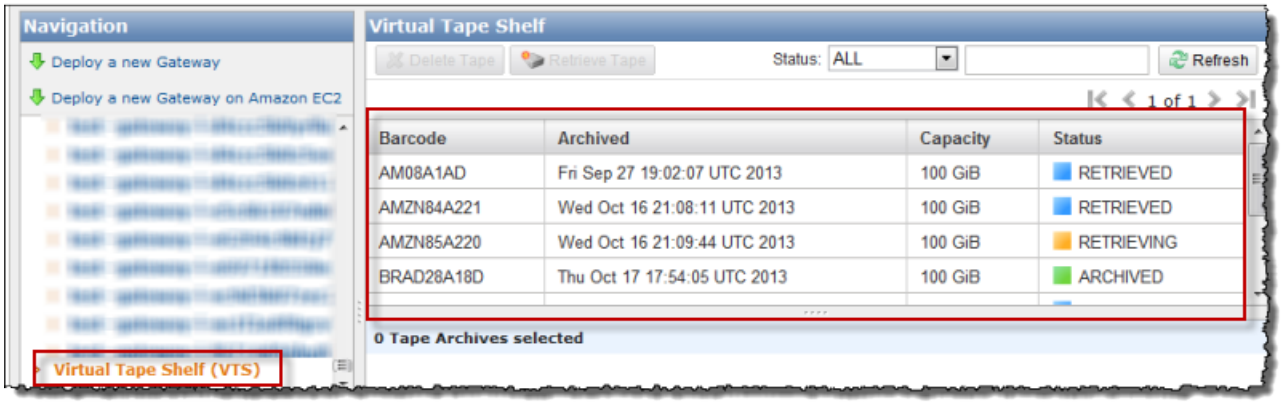


Each virtual tape in the VTS has an associated status. For more information, see Tape Status Information (p. 286). Over time, you might collect a large number of tapes on the shelf, so the console

supports search for a tape by barcode, status, or both. When you search by barcode, you can filter by status.

### To search by barcode and status

1. In the AWS Storage Gateway console navigation panel, click **Virtual Tape Shelf**, and then type the barcode of the virtual tape you want to find in the **Search Barcode** box. By default, AWS Storage Gateway searches for all virtual tapes. However, you can also filter your search by status.
2. In the **Status** drop-down list box, select the status of the tapes you want to find. The tapes that match your criteria appear in a grid.

### To search by status

1. In the AWS Storage Gateway console, click **Virtual Tape Shelf**, and then select the status to search by from the **Status** drop-down list box. All virtual tapes that have the status that you specify will appear in the VTS.

   For information about tape status in the VTS, see Tape Status Information (p. 286).
2. If the virtual tapes are not found, click **Refresh** to locate them.

## Tape Status Information

As shown in the preceding console screenshot, each tape in VTS has an associated status.

### To determine the status of a virtual tape

1. In the AWS Storage Gateway console, click **Virtual Tape Shelf**.
2. In the **Status** column of the grid, check the status of the tape.

   The tape status also appears in the **Details** tab of each virtual tape.

The following table lists and describes the possible status values.

| Virtual Tape Status in VTS | Description |
|---|---|
| ARCHIVED | The virtual tape has been ejected and is uploaded to the VTS. |
| RETRIEVING | The virtual tape is being retrieved from the VTS.<br>**Note**<br>The virtual tape can be retrieved only to a gateway-VTL. |
| RETRIEVED | The virtual tape is retrieved from the VTS. The retrieved tape is read-only. |

## Retrieving Tapes from the VTS to Your Gateway

To access virtual tape data in the VTL, you must first retrieve the virtual tape you want from a VTS to your gateway-VTL. AWS Storage Gateway provides one VTL per gateway, but there is only one VTS per

account per region. You can restore a virtual tape from VTL in a region to a gateway-VTL in the same region.

If you have multiple gateway-VTLs deployed in a region, you can retrieve a tape to only one gateway.

The retrieved tape is write-protected; you can only read the data on the tape.

**Important**
It takes up to 24 hours for the tape to be available in your gateway-VTL.

**Note**
There is a charge for retrieving tapes from the VTS. For detailed pricing information, see AWS Storage Gateway Pricing.

### To retrieve a tape to your gateway

1. In the navigation pane of the AWS Storage Gateway console, click **Virtual Tape Shelf**. The console displays all virtual tapes that have been archived by all your gateways.
2. Select the virtual tape you want to retrieve, and click **Retrieve Tape**.

   **Note**
   The status of the virtual tape you want to retrieve must be ARCHIVED.

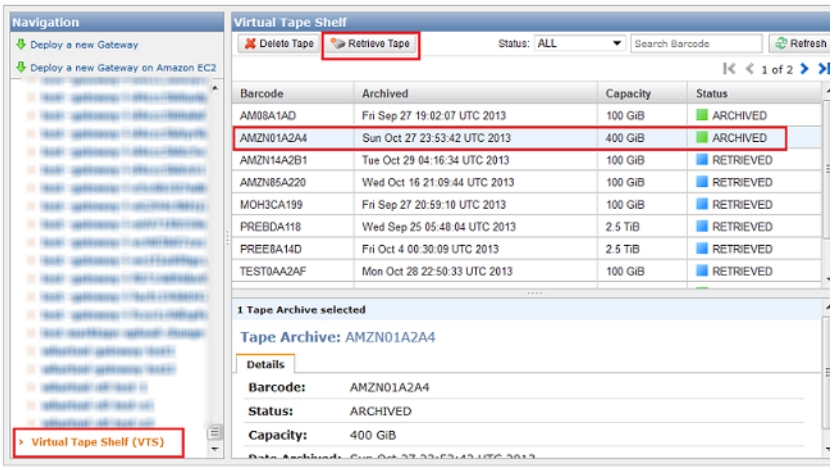


3. In the **Tape Barcode** box of the **Retrieve Tape** wizard, verify that the barcode identifies the virtual tape you want to retrieve.

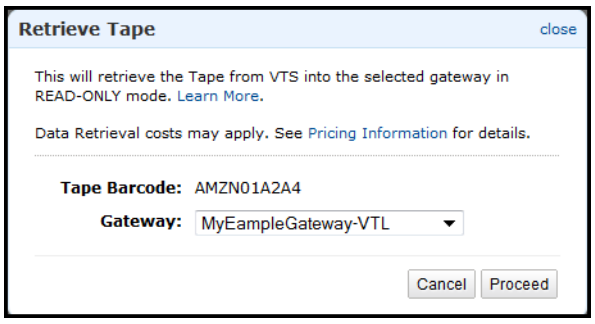
Retrieve Tape
close
This will retrieve the Tape from VTS into the selected gateway in READ-ONLY mode. Learn More.
Data Retrieval costs may apply. See Pricing Information for details.
Tape Barcode: AMZN01A2A4
Gateway: MyEampleGateway-VTL
Cancel
Proceed

4. In the **Gateway** drop-down list, select the gateway you want to retrieve the archived tape to, and then click **Proceed**.

The status of the tape changes from ARCHIVED to RETRIEVING. At this point, your data is being moved from the virtual tape shelf (backed by Amazon Glacier) to the virtual tape library (backed by Amazon S3). After all the data is moved, the status of the virtual tape in the VTS changes to RETRIEVED.

**Note**
Retrieved virtual tapes are read-only.

## Deleting Tapes from the VTS

If you no longer need a virtual tape, you can delete it from the virtual tape shelf.

**Note**
There is a charge for early deletion of tapes from the VTS. For detailed pricing information, see AWS Storage Gateway Pricing.

### To delete a virtual tape from the VTS

1. In the AWS Storage Gateway console, click **Virtual Tape Shelf**.
2. Select the virtual tape you want to delete, and click the **Delete Tape** button.

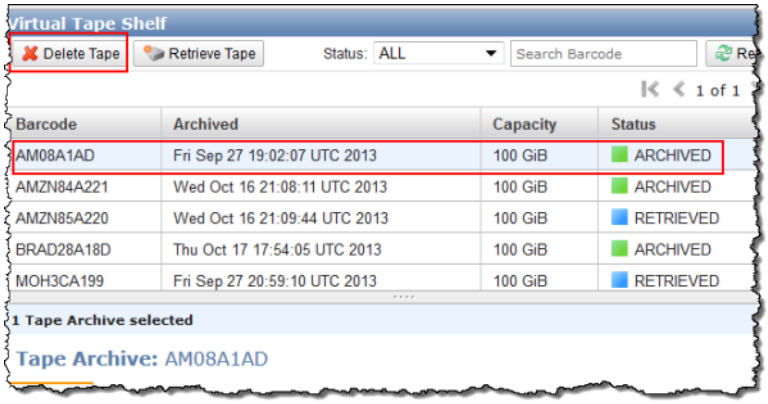


If a virtual tape has RETRIEVING status, the retrieval must be canceled before deleting the tape. For more information, see Canceling Tape Retrieval (p. 281). If the tape has RETRIEVED status, it must be archived before you delete it. For more information, see Archiving Tapes to the VTS (p. 281).

# Connecting to the VTL Devices on Your Gateway

**Topics**


- Customizing Windows iSCSI Settings (p. 289)
- Connecting from a Red Hat Client to Your VTL Devices (p. 291)
- Configuring CHAP Authentication for Your VTL Devices (p. 294)


The gateway-VTL exposes several tape drives and a media changer, referred to collectively as VTL devices, as Internet Small Computer System Interface (iSCSI) targets. For more information, see Requirements (p. 11).

**Note**
You connect only one application to each iSCSI target.

The following diagram highlights the iSCSI target in the larger picture of the AWS Storage Gateway architecture. For more information on AWS Storage Gateway architecture, see Gateway–Virtual Tape Library (Gateway-VTL) Architecture (p. 5).

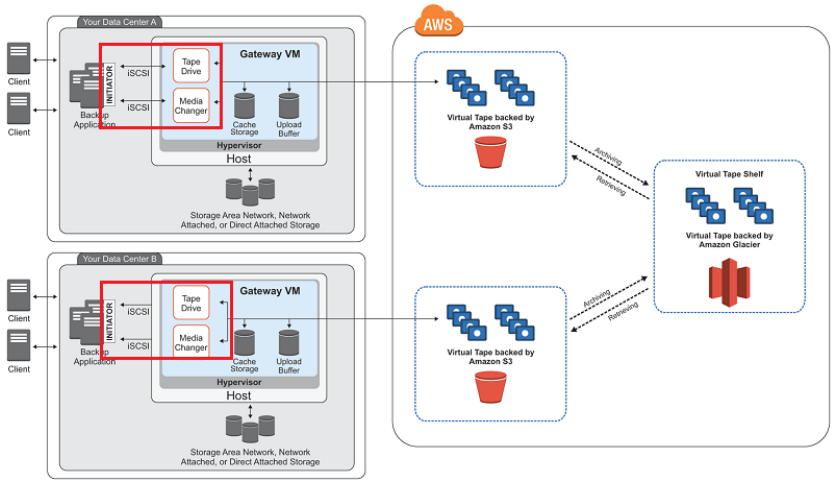


The iSCSI standard is an Internet Protocol (IP)–based storage networking standard for initiating and managing connections between IP-based storage devices and clients.

The following list defines some of the iSCSI terms that are used to describe the connection and the components involved.

iSCSI initiator
: The client component of an iSCSI network. The initiator sends requests to the iSCSI target. Initiators can be implemented in software or hardware. AWS Storage Gateway only supports software initiators.

iSCSI target
: The server component of the iSCSI network that receives and responds to requests from initiators. Each of your storage volumes is exposed as an iSCSI target. Connect only one iSCSI initiator to each iSCSI target.

Microsoft iSCSI initiator
: The software program on Windows computers that enables you to connect a client computer (that is, the computer running the application whose data you want to write to the gateway) to an external iSCSI-based array (that is, the gateway). The connection is made using the host computer's Ethernet network adapter card. The Microsoft iSCSI initiator is already installed on Windows Server 2008 R2, Windows 7, Windows Server 2008, and Windows Vista.

Red Hat iSCSI initiator
: The iscsi-initiator-utils Resource Package Manager (RPM) package provides you with an iSCSI initiator implemented in software for Red Hat. The package includes a server daemon for the iSCSI protocol.

You can optionally configure Challenge-Handshake Authentication Protocol (CHAP) for either client type.

## Customizing Windows iSCSI Settings

When using a Windows client, you use the Microsoft iSCSI initiator to connect to your VTL devices. This is a two-step process:

1. Connect gateway-VTL devices to your Windows client.
2. If you are using a backup application, configure the application to use the devices.

The Getting Started example setup provides instructions for both these steps. It uses the Symantec NetBackup backup application. For more information, see Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245) and Configuring NetBackup Storage Devices (p. 250).

## Recommended Windows iSCSI Settings

After setting up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets. By increasing the iSCSI timeout values as shown in the following steps, you improve the ability of your application to deal with write operations that take a long time and other transient issues such as network interruptions.

**Note**
Before making changes to the registry, you should make a backup copy of the registry. For information on making a backup copy and other best practices to follow when working with the registry, see Registry best practices in the *Windows Server TechCenter*.

### To customize your Windows iSCSI settings

1. Increase the maximum time for which requests are queued.

   a. Start Registry Editor (`Regedit.exe`).

   b. Navigate to the globally unique identifier (GUID) key for the device class that contains iSCSI controller settings, shown following.

   **Warning**
   Make sure you are working in the **CurrentControlSet** subkey and not another control set like **ControlSet001** or **ControlSet002**.

   ```
   HK_Local_Machine\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}
   ```

   c. Find the subkey for the Microsoft iSCSI initiator.

   The key will be represented by a four-digit number such as `0000` or `00001`.

   ```
   HK_Local_Machine\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>
   ```

   Depending on what is installed on your computer, the Microsoft iSCSI initiator might not be the subkey `0000`. You can ensure that you have selected the correct subkey by verifying that the string `DriverDesc` has the value `Microsoft iSCSI Initiator` as shown in the following example.

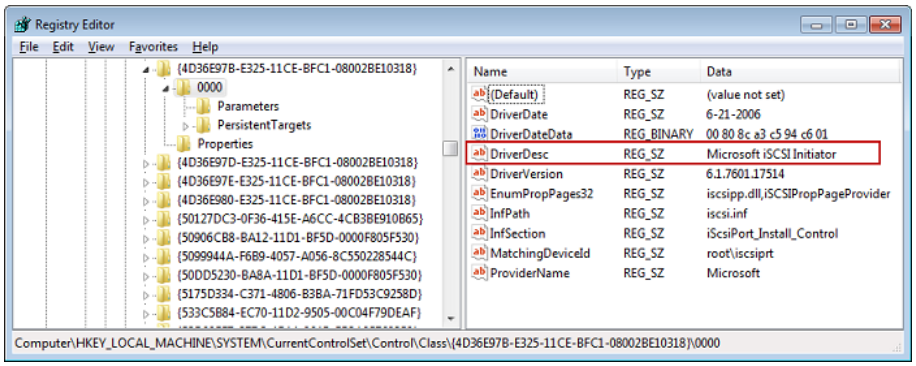


   d. Click the **Parameters** subkey to show the iSCSI settings.

   e. Right-click the **MaxRequestHoldTime** DWORD (32-bit) value, right-click **Modify**, and change its value to 600.

This value represents a hold time of 600 seconds. The example following shows the **MaxRequestHoldTime** DWORD value with a value of 600.

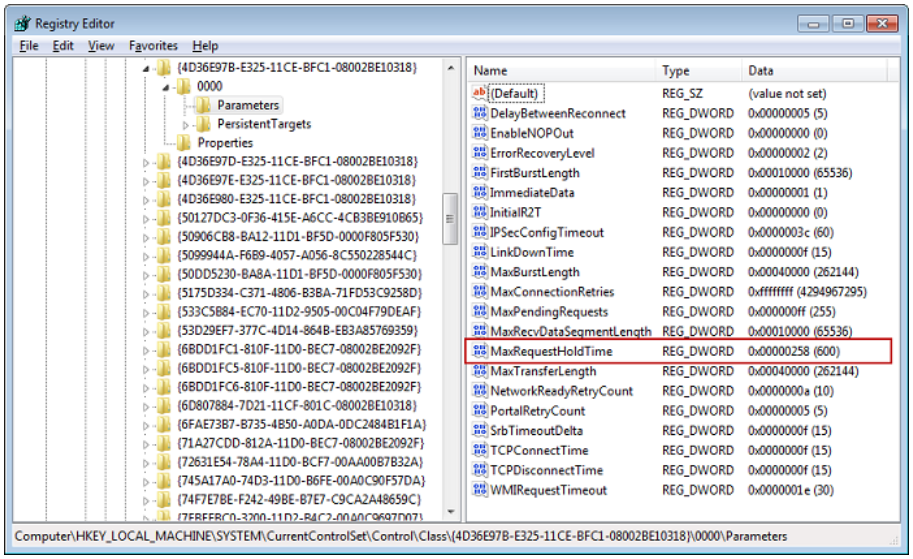


2. Increase the disk timeout value.

   a. Start Registry Editor (`Regedit.exe`).

   b. Navigate to the **Disk** subkey in the **Services** subkey of the **CurrentControlSet**, shown following.

   ```
   HK_Local_Machine\SYSTEM\CurrentControlSet\Services\Disk
   ```

   c. Right-click the **TimeoutValue** DWORD (32-bit) value, click **Modify**, and change its value to 600.

   This value represents a timeout of 600 seconds.

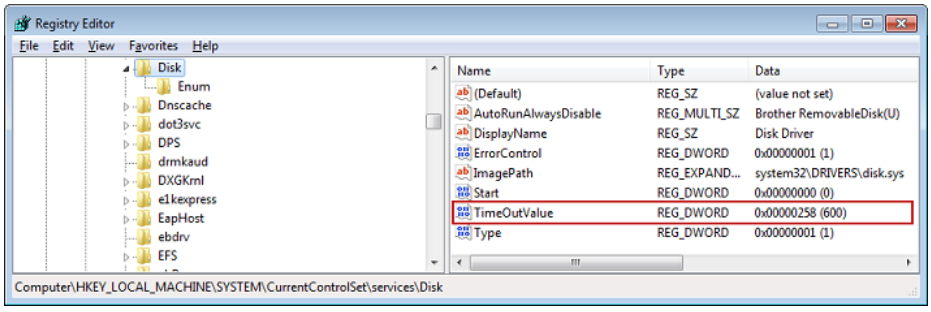


3. Restart your system to ensure that the new configuration values take effect.

   Before restarting, you must make sure that the results of all write operations to virtual tapes are flushed. To do this, take any mapped virtual tapes offline before restarting.

## Connecting from a Red Hat Client to Your VTL Devices

When using Red Hat Linux, you use the iscsi-initiatorutils RPM package to connect to your VTL devices.

**Note**
After completing the following procedure, if you are using backup software, you will also need to configure the backup application to use the VTL devices.

### To connect a Linux client to VTL devices

1. Install the iscsi-initiator-utils RPM package, if it isn't already installed on your client.

   You can use the following command to install the package.

   ```
   sudo yum install iscsi-initiator-utils
   ```

2. Ensure that the iSCSI daemon is running.

   a. Verify that the iSCSI daemon is running using the following command.

      ```
      sudo /etc/init.d/iscsi status
      ```

   b. If the status command does not return a status of *running*, then start the daemon using the following command.

      ```
      sudo /etc/init.d/iscsi start
      ```

3. To discover the VTL device targets defined for a gateway-VTL, use the following discovery command.

   ```
   sudo /sbin/iscsiadm --mode discovery --type sendtargets --portal GATE
   WAY_IP:3260
   ```

   Substitute your gateway's IP address for the *GATEWAY_IP* variable in the preceding command. You can find the gateway IP in the **iSCSI Target Info** properties of a VTL device on the AWS Storage Gateway console.

   The output of the discovery command will look like the following example output.

   ```
   iqn.1997-05.com.amazon:gatewayId-tapedrive-01
   ```

   Your iSCSI qualified name (IQN) will be different than what is shown preceding, because IQN values are unique to an organization. When you activated your gateway-VTL, 10 virtual tape drives and one media changer were created for you as targets. When you select a virtual tape drive or media changer on the AWS Storage Gateway console, you can find the target name of that device in the **iSCSI Target Info** properties pane.

4. To connect to a target, use the following command.

   You need to specify the correct *GATEWAY_IP* and IQN in the connect command.

   **Warning**
   For gateway-VTLs that are deployed on an Amazon EC2 instance, accessing the gateway over a public Internet connection is not supported. The elastic IP address of the Amazon EC2 instance cannot be used as the target address.

   ```
   sudo /sbin/iscsiadm --mode node --targetname iqn.1997-05.com.amazon:mydevice
    --portal GATEWAY_IP:3260,1 --login
   ```

5. To verify that the devices are attached to the client machine (the initiator), use the following command.

```
ls -l /dev/disk/by-path
```

After setting up your initiator, we highly recommend that you customize your iSCSI settings as discussed in Customizing Your Linux iSCSI Settings (p. 293).

## Customizing Your Linux iSCSI Settings

After setting up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets. By increasing the iSCSI timeout values as shown following, you improve the ability of your application to deal with write operations that take a long time and other transient issues such as network interruptions.

### To customize your Linux iSCSI settings

1. Increase the maximum time for which requests are queued.

   a. Open the `/etc/iscsi/iscsid.conf` file and find the following lines.

```
node.session.timeo.replacement_timeout = [replacement_timeout_value]
node.conn[0].timeo.noop_out_interval = [noop_out_interval_value]
node.conn[0].timeo.noop_out_timeout = [noop_out_timeout_value]
```

   b. Set the *[replacement_timeout_value]* value to 600.

   Set the *[noop_out_interval_value]* value to 60.

   Set the *[noop_out_timeout_value]* value to 600.

   All three values are in seconds.

   **Note**
   The `iscsid.conf` settings must be made before discovering the gateway. If you have already discovered your gateway and/or logged in to the target, you can delete the entry from the discovery database using the following command. Then you can rediscover or log in again to pick up the new configuration.

```
iscsiadm -m discoverydb -t sendtargets -p gateway_ip:3260 -o delete
```

2. Increase the disk timeout value in the rules file.

   a. If you are using the Red Hat Enterprise Linux 5 initiator, open the `/etc/udev/rules.d/50-udev.rules` file and find the following line.

```
ACTION=="add", SUBSYSTEM=="scsi" , SYSFS{type}=="0|7|14", \
RUN+="/bin/sh -c 'echo [timeout] > /sys$$DEVPATH/timeout'"
```

   **Note**
   This rules file does not exist in Red Hat Enterprise Linux 6 initiators, so you must create it.

To create the rules file in Red Hat Enterprise Linux 6, use the following command and then add the lines of code as shown above.

```
sudo vim /etc/udev/rules.d/50-udev.rules
```

b. Set the *timeout* value to 600.

This value represents a timeout of 600 seconds.

3. Restart your system to ensure that the new configuration values take effect.

Before restarting, you must make sure that the results of all write operations to your storage volumes are flushed. To do this, unmount storage volumes before restarting.

4. You can test the configuration by using the following command.

```
udevadm test PATH_TO_ISCSI_DEVICE
```

This command shows the udev rules that are applied to the iSCSI device.

# Configuring CHAP Authentication for Your VTL Devices

AWS Storage Gateway supports authentication between your gateway and iSCSI initiator via Challenge-Handshake Authentication Protocol (CHAP). CHAP provides protection against playback attacks by periodically verifying the identity of an iSCSI initiator as authenticated to access a VTL device target. To set up CHAP, you must configure it both on the AWS Storage Gateway console and in the iSCSI initiator software that you use to connect to the target. This section discusses mutual CHAP, which is when the initiator authenticates the target and the target authenticates the initiator.

**Note**
Note that CHAP must be configured separately for each VTL device.

### To set up mutual CHAP for your gateway-VTL

1. Configure CHAP on the AWS Storage Gateway console, as discussed in To configure CHAP on the AWS Storage Gateway console (p. 117).
2. In your client initiator software, complete the CHAP configuration.
   - To configure mutual CHAP on a Windows client, see To configure mutual CHAP on a Windows client (p. 118).
   - To configure mutual CHAP on a Red Hat Linux client, see To configure mutual CHAP on a Red Hat Linux client (p. 123).

### To configure CHAP in the AWS Storage Gateway console

In this procedure, you specify two secret keys that are used to read and write to a virtual tape in a virtual tape drive. These same keys are used in the procedure to configure the client initiator.

1. On the AWS Storage Gateway console, go to the **iSCSI Target Info** tab of the virtual tape drive for which you want to configure CHAP.
2. Click the **Configure CHAP Authentication** link.

In this screenshot, the device selected for CHAP configuration is virtual tape drive-01.

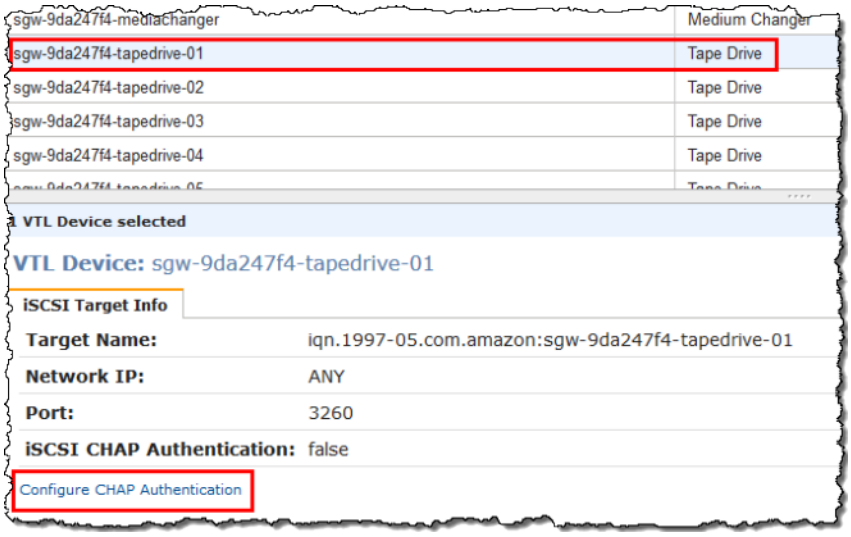


3. Configure CHAP in the **Configure CHAP Authentication** dialog box.

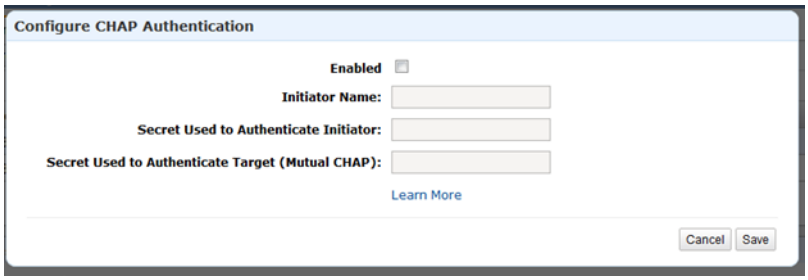


   a. Check the **Enabled** box.
   b. Type the name of your iSCSI initiator in the **Initiator Name** box.

      The initiator name can be found using your iSCSI initiator software. For example, for Windows clients, the name is the value in the **Configuration** tab of the iSCSI initiator. For more information, see To configure mutual CHAP on a Windows client (p. 296).

      **Note**
      To change an initiator name, you must first disable CHAP, change the initiator name in your iSCSI initiator software, and then enable CHAP with the new name.

   c. Type the secret requested in the **Secret used to Authenticate Initiator** box.

      This secret must be at least 12 characters long. It is the secret key that the initiator (that is, Windows client) must know to participate in CHAP with the target.

   d. Type the secret requested in the **Secret used to Authenticate Target (Mutual CHAP)** box.

      This secret must be at least 12 characters long. It is the secret key that the target must know to participate in CHAP with the initiator.

      **Note**
      The secret used to authenticate the target must be different than the secret to authenticate the initiator.

   e. Click **Save**.
   f. In the confirmation dialog box, click **Close**. The **iSCSI Target Info** tab now indicates that CHAP authentication is being used and is enabled.

4. In the **iSCSI Target Info** tab, verify that the **iSCSI CHAP Authentication** is set to *true*.

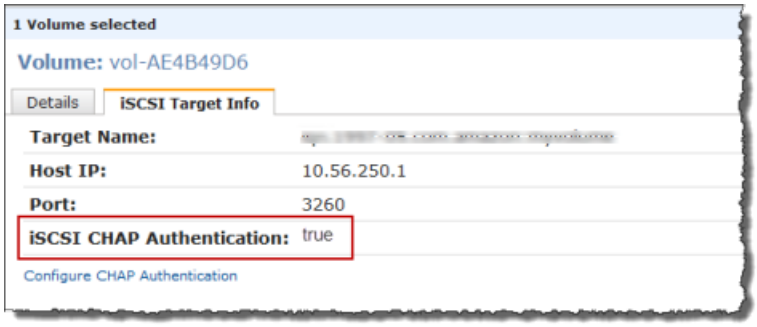


### To configure mutual CHAP on a Windows client

In this procedure, you configure CHAP in the Microsoft iSCSI initiator using the same keys that you used to configure CHAP for the virtual tape drive in the console.

1. If the iSCSI initiator is not already started, on the **Start** menu of your Windows client computer, click **Run**, type **`iscsicpl.exe`**, and then click **OK** run the program.
2. Configure mutual CHAP configuration for the initiator (that is, the Windows client).

   a. Click the **Configuration** tab.

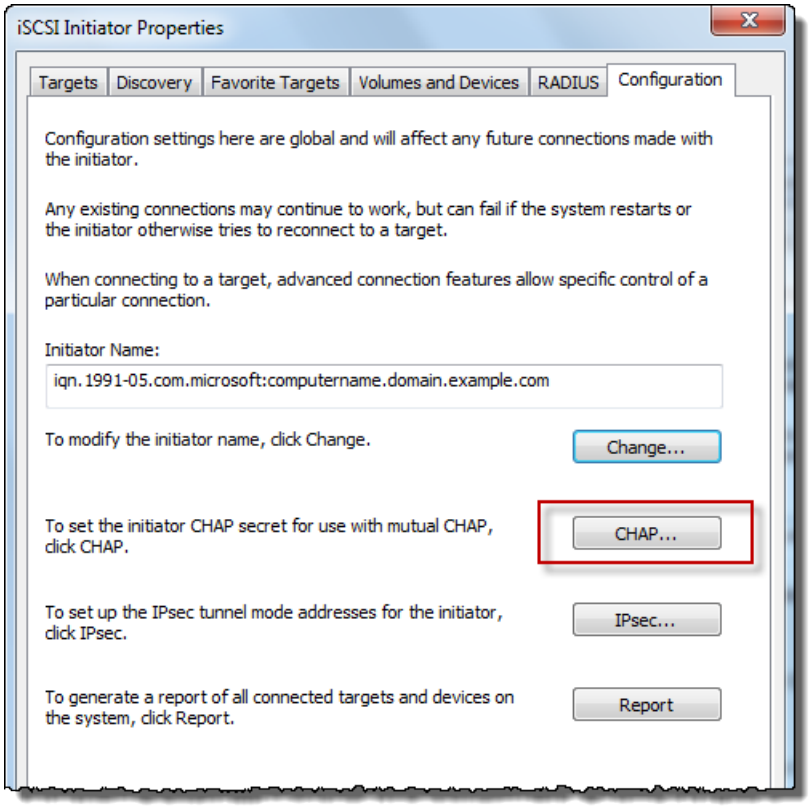
iSCSI Initiator Properties
Targets
Discovery
Favorite Targets
Volumes and Devices
RADIUS
Configuration
Configuration settings here are global and will affect any future connections made with the initiator.
Any existing connections may continue to work, but can fail if the system restarts or the initiator otherwise tries to reconnect to a target.
When connecting to a target, advanced connection features allow specific control of a particular connection.
Initiator Name:
iqn.1991-05.com.microsoft:computername.domain.example.com
To modify the initiator name, click Change.
Change...
To set the initiator CHAP secret for use with mutual CHAP, click CHAP.
CHAP...
To set up the IPsec tunnel mode addresses for the initiator, click IPsec.
IPsec...
To generate a report of all connected targets and devices on the system, click Report.
Report

**Note**

Note that the **Initiator Name** box will be unique to your initiator and company. The name shown here is the value that you used in the **Configure CHAP Authentication** dialog box of the AWS Storage Gateway console.
The name shown in the example image is for demonstration purposes only.

b. Click **CHAP**.

c. In the **iSCSI Initiator Mutual Chap Secret** dialog box, type the mutual CHAP secret value.

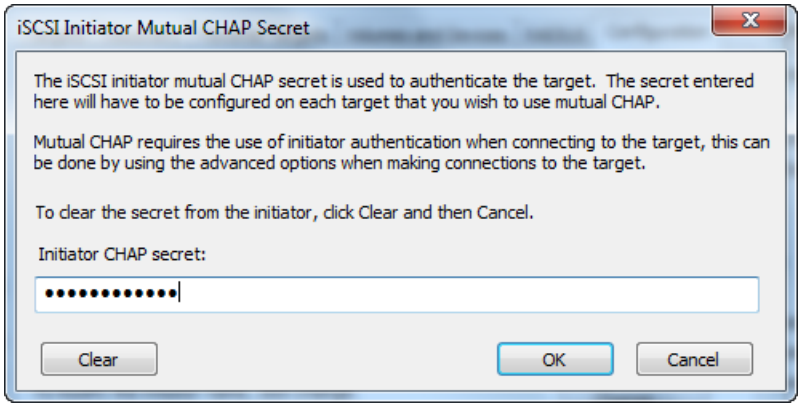
iSCSI Initiator Mutual CHAP Secret
The iSCSI initiator mutual CHAP secret is used to authenticate the target. The secret entered here will have to be configured on each target that you wish to use mutual CHAP.
Mutual CHAP requires the use of initiator authentication when connecting to the target, this can be done by using the advanced options when making connections to the target.
To clear the secret from the initiator, click Clear and then Cancel.
Initiator CHAP secret:
Clear
OK
Cancel

In this dialog box, you are entering the secret that the initiator (the Windows client) uses to authenticate the target (the VTL device). This secret allows the target to read and write to the initiator. This secret maps to the secret typed into the **Secret used to Authenticate Target (Mutual CHAP)** box in the **Configure CHAP Authentication** dialog box. For more information, see Configuring CHAP Authentication for Your VTL Devices (p. 294).

d. If the key that you typed is less than 12 characters, an **Initiator CHAP secret** error dialog box appears.

Click **OK**, and then type the key again.

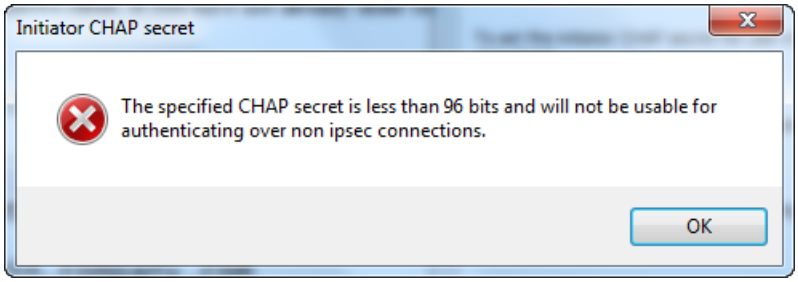
Initiator CHAP secret
The specified CHAP secret is less than 96 bits and will not be usable for authenticating over non ipsec connections.
OK

3. Configure the target with the initiator's secret to complete the mutual CHAP configuration.

a. Click the **Targets** tab.

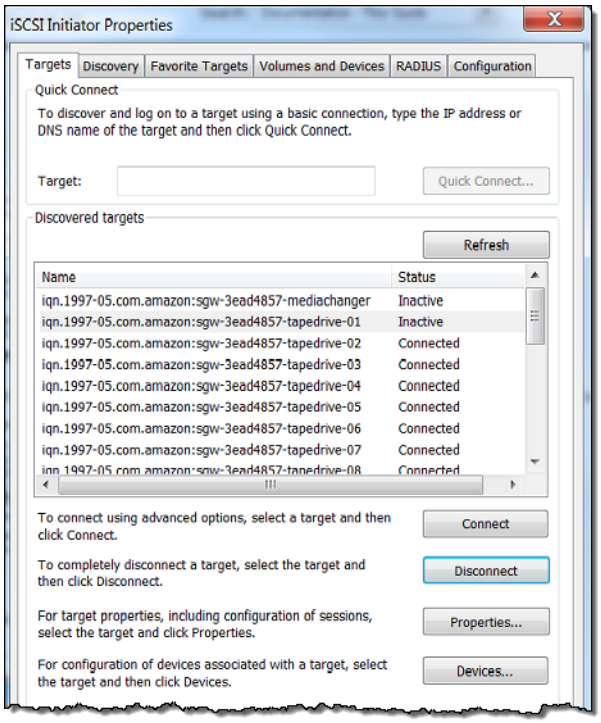


b. If the target that you want to configure for CHAP is currently connected, disconnect the target by selecting it and clicking **Disconnect**.

c. Select the target that you want to configure for CHAP, and then click **Connect**.

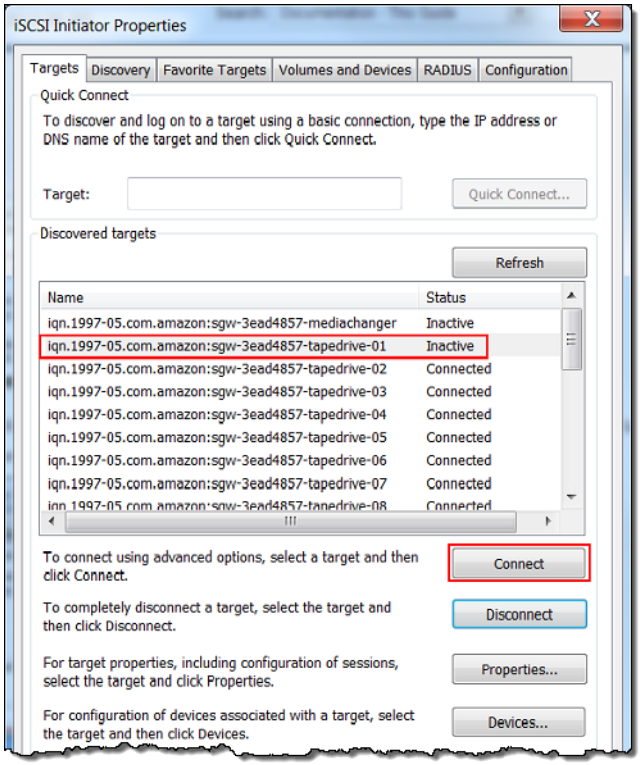


d. In the **Connect to Target** dialog box, click **Advanced**.

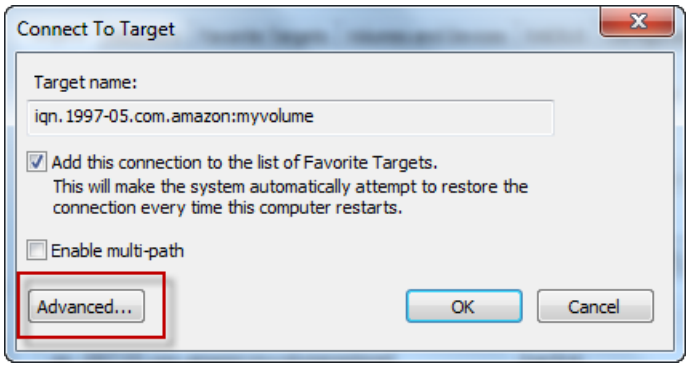
Connect To Target
Target name:
iqn.1997-05.com.amazon:myvolume
Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.
Enable multi-path
Advanced...
OK
Cancel

e. In the **Advanced Settings** dialog box, configure CHAP.

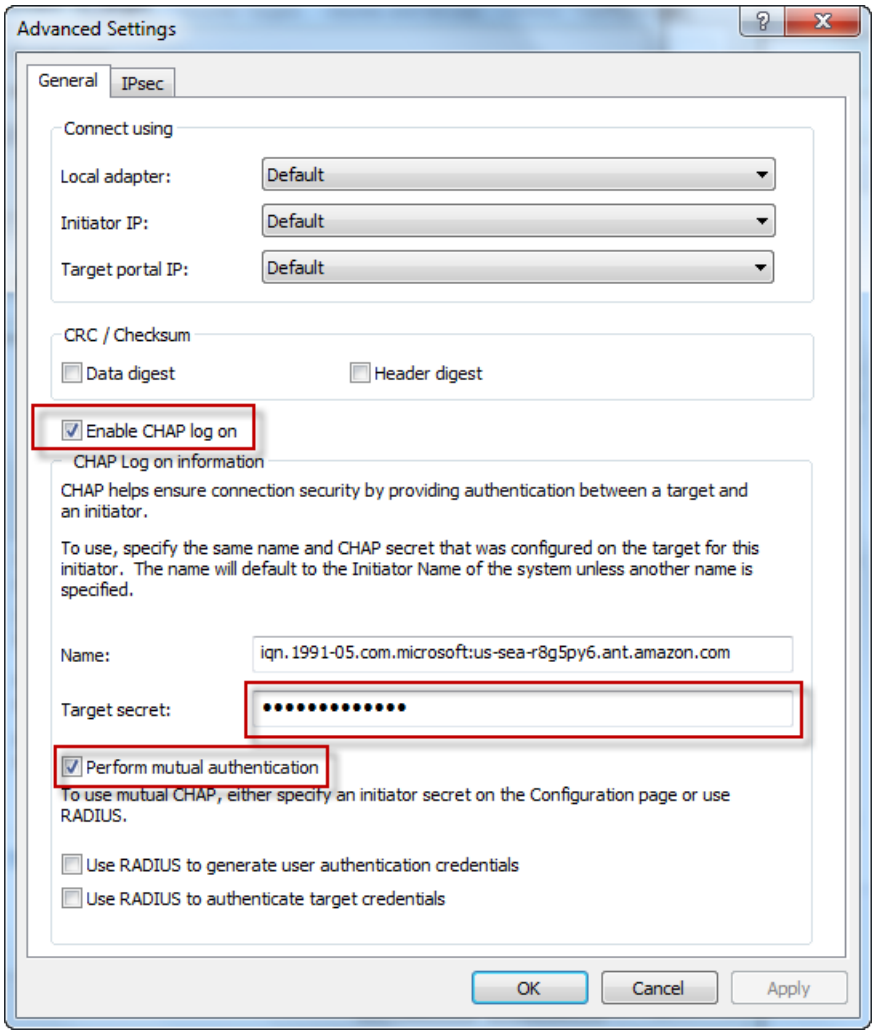
Advanced Settings
General
IPsec
Connect using
Local adapter:
Default
Initiator IP:
Default
Target portal IP:
Default
CRC / Checksum
Data digest
Header digest
Enable CHAP log on
CHAP Log on information
CHAP helps ensure connection security by providing authentication between a target and an initiator.
To use, specify the same name and CHAP secret that was configured on the target for this initiator. The name will default to the Initiator Name of the system unless another name is specified.
Name:
iqn.1991-05.com.microsoft:us-sea-r8g5py6.ant.amazon.com
Target secret:
Perform mutual authentication
To use mutual CHAP, either specify an initiator secret on the Configuration page or use RADIUS.
Use RADIUS to generate user authentication credentials
Use RADIUS to authenticate target credentials
OK
Cancel
Apply

i. Click to select the **Enable CHAP log on** check box.

ii. Type the secret that is required to authenticate the initiator. This secret maps to the secret typed into the **Secret used to Authenticate Initiator** box in the **Configure CHAP Authentication** dialog box. For more information, see Configuring CHAP Authentication for Your VTL Devices (p. 294).

iii. Click to select the **Perform mutual authentication** check box.

iv. Click **OK** to apply the changes.

f. In the **Connect to Target** dialog box, click **OK**.

4. If you provided the correct secret key, the target shows a status of **Connected**.

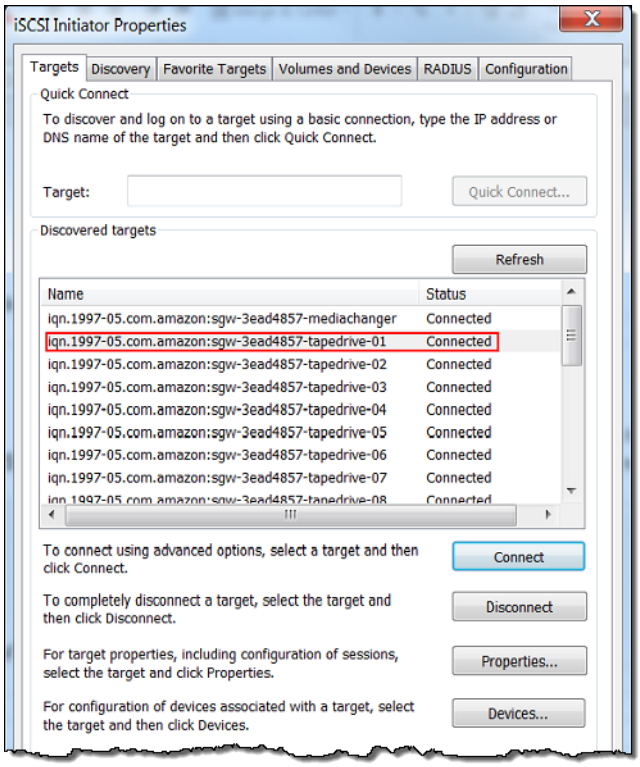


The following procedure assumes that the iSCSI daemon is running and that you have already connected to a target. If you have not completed these two tasks, see Connecting from a Red Hat Client to Your VTL Devices (p. 291).

### To configure mutual CHAP on a Red Hat Linux client

In this procedure, you configure CHAP in the Linux iSCSI initiator using the same keys that you used to configure CHAP for the virtual tape drive in the console.

1. Disconnect and remove any existing configuration for the target for which you are about to configure CHAP.

   a. To find the target name and ensure it is a defined configuration, list the saved configurations using the following command.

   ```
   sudo /sbin/iscsiadm --mode node
   ```

   b. Disconnect from the target.

The following command disconnects from the target named *mydevice* that is defined in the Amazon iSCSI qualified name (IQN). Change the target name and IQN as required for your situation. Also, substitute your gateway-VTL's IP address for the *GATEWAY_IP* variable.

```
sudo /sbin/iscsiadm --mode node --logout GATEWAY_IP:3260,1 iqn.1997-
05.com.amazon:mydevice
```

c. Remove the configuration for the target.

The following command removes the configuration for the *mydevice* target.

```
sudo /sbin/iscsiadm --mode node --op delete --targetname iqn.1997-
05.com.amazon:mydevice
```

2. Edit the iSCSI configuration file to enable CHAP.

a. Get the name of the initiator (the client you are using).

The following command gets the initiator name from the `/etc/iscsi/initiatorname.iscsi` file.

```
sudo cat /etc/iscsi/initiatorname.iscsi
```

The output from this command will look like this:

```
InitiatorName=iqn.1994-05.com.redhat:8e89b27b5b8
```

b. Open the `/etc/iscsi/iscsid.conf` file.

c. Uncomment the following lines in the file and specify the correct username and passwords (that is, the secret keys).

```
node.session.auth.authmethod = CHAP
node.session.auth.username = username
node.session.auth.password = password
node.session.auth.username_in = username_in
node.session.auth.password_in = password_in
```

Fill in the preceding items using the following table as guidance.

| Configuration Setting | Value |
|---|---|
| `username` | The initiator name that you found in a previous step in this procedure. The value will start with *iqn*. For example, **`iqn.1994-05.com.redhat:8e89b27b5b8`** is a valid `username`. |
| `password` | The secret key used to authenticate the initiator (the client you are using) when it communicates with the VTL device. |
| `username_in` | The IQN of the target VTL device. The value will start with *iqn* and end with the target name. For example, **`iqn.1997-05.com.amazon:mydevice`** is a valid `username_in`. |

| Configuration Setting | Value |
|---|---|
| `password_in` | The secret key used to authenticate the target (the VTL device) when it communicates to the initiator. |

d. Save the changes in the configuration file, and then close the file.

3. Discover and log in to the target.

   You can follow the steps provided in Connecting from a Red Hat Client to Your VTL Devices (p. 291) to do so.

# Performing Gateway-VTL Maintenance Tasks

AWS Storage Gateway supports several maintenance tasks after you activate your gateway. This section groups these tasks as follows:

- **Tasks you perform in the AWS Storage Gateway console –** For more information, see Common Gateway-VTL Maintenance Tasks (p. 302).
- **Tasks you perform on-premises –** If your gateway-VTL is deployed on-premises, you perform most of the maintenance tasks in the local console of your gateway. For more information, see Maintenance Tasks When Running a Gateway-VTL On-Premises (p. 321).
- **Tasks you perform when the gateway-VTL is running as an Amazon EC2 instance –** For more information, see Maintenance Tasks When Running a Gateway-VTL on an Amazon EC2 Instance (p. 342).

## Common Gateway-VTL Maintenance Tasks

This section describes gateway-VTL maintenance tasks that you can perform in the AWS Storage Gateway console. For maintenance tasks specific to a gateway-VTL running on an Amazon EC2 instance, see Maintenance Tasks When Running a Gateway-VTL on an Amazon EC2 Instance (p. 342).

**Topics**

- Managing the Upload Buffer and Cache Storage (Gateway-VTL) (p. 304)
- Shutting Down and Starting Your Gateway-VTL (p. 311)
- Managing Gateway-VTL Updates Using the AWS Storage Gateway Console (p. 312)
- Updating Bandwidth Rate Limits for Your Gateway-VTL (p. 313)
- Disabling Your Gateway-VTL (p. 319)
- Deleting Your Gateway-VTL (p. 320)

You can perform many tasks related to gateway maintenance on the **Gateway** tab in the AWS Storage Gateway console. The following screenshot shows the **Gateway** tab.

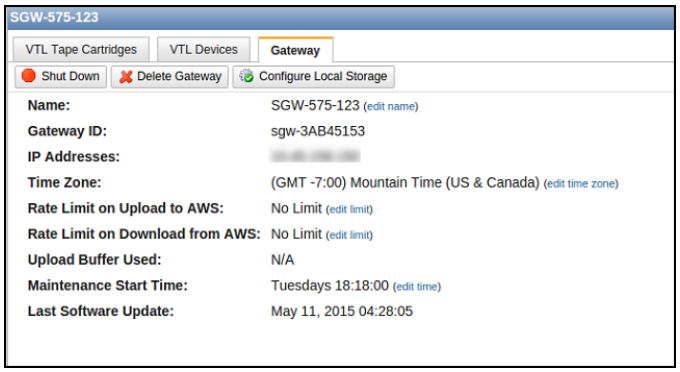


The following table summarizes the information on the **Gateway** tab that you can update. To edit the value for a particular item, click the **edit** link at the end of the box for that item.

| Maintenance Item | Comments |
|---|---|
| **Name** | The gateway name. You can change this name as needed. If you use Amazon CloudWatch to view your gateway metrics (see Using Amazon CloudWatch Metrics (p. 352)), you might want to take note of the previous name and the new name to avoid confusion. You can also locate your gateway using the gateway ID, which remains the same through the life of the gateway. |
| **Gateway ID** | The gateway identifier. AWS Storage Gateway assigns a unique identifier for each gateway. This value cannot be changed. |
| **IP Addresses** | The Internet Protocol (IP) addresses you use when connecting your initiator to the VTL devices on your gateway. If you can configure your gateway to use multiple network interface cards (NICs), then all those IP addresses appear here. |
| **Time Zone** | The time zone. AWS Storage Gateway uses this item when displaying time-based information. For gateway-VTL setups, the time zone specified is used in setting the maintenance start time. |
| **Rate Limit on Upload to AWS** | The limit set for upload throughput from the gateway to AWS. Using this item, you can apply bandwidth throttling to your gateway to control the amount of network bandwidth used for uploads. Specify the rate limit as kilobits per second (kbps). The default is no rate limit on upload. For more information about updating this bandwidth, see Updating Bandwidth Rate Limits for Your Gateway-VTL (p. 313). |
| **Rate Limit on Download from AWS** | The limit set for download throughput from AWS to your gateway. Using this item, you can apply bandwidth throttling to your gateway to control the amount of network bandwidth used for downloads. Specify the rate limit as kilobits per second (kbps). The default is no rate limit on download. For more information about updating this bandwidth, see Updating Bandwidth Rate Limits for Your Gateway-VTL (p. 313). Bandwidth throttling can affect performance when reading data from virtual tapes. |
| **Upload Buffer Used** | The upload buffer used. This value cannot be changed. For information about how to monitor the upload buffer and how it changes over time, see Monitoring the Upload Buffer (p. 355). |

| Maintenance Item | Comments |
|---|---|
| **Maintenance Start Time** | The beginning of the maintenance time period. Each gateway has a maintenance window of one time period per week. During activation, a default time period is assigned to your gateway. To change that maintenance time, click **edit** for this item, and then specify a day of the week and time of the day in the time zone of the gateway. |
| **Last Software Update** | The date and time the last successful software update was applied to the gateway. |
| **Apply Update Now** | A message button. If there is an update for your AWS Storage Gateway installation, a message box displaying this button appears on the console. To apply the update immediately, click **Apply Update Now**. If you do not apply the update immediately, AWS Storage Gateway applies the update based on your **Maintenance Start Time** setting. For more information, see Managing Gateway-VTL Updates Using the AWS Storage Gateway Console (p. 312). |

## Managing the Upload Buffer and Cache Storage (Gateway-VTL)

After the initial configuration and deployment of your gateway-VTL, you might find that you need to adjust the local storage by adding or removing disks for an upload buffer or adding disks for cache storage.

> **Important**
> AWS Storage Gateway does not support removing disks allocated for cache storage.

The amount of upload buffer space your gateway-VTL requires depends on several factors, such as the rate of incoming data to the virtual tapes, the rate of outgoing data to AWS, and your network bandwidth. If your applications continue to write data at a fast rate to your virtual tapes, and network throughput is not sufficient for the gateway-VTL to upload data to AWS, then eventually your upload buffer will be filled with data waiting to be uploaded to AWS. Your gateway's ability to periodically create recovery points of your tapes depends on having a properly configured upload buffer. If your upload buffer is too small, it might fill up while writing to a tape, which could affect how quickly the gateway-VTL is able to create the next recovery point. For information about sizing the upload buffer, see Sizing the Upload Buffer (p. 305). For more information about recovery points, see Recovering a Virtual Tape from a Malfunctioning Gateway-VTL (p. 282).

The amount of cache storage capacity should be at least the size of the upload buffer. This guideline helps ensure that the cache storage is large enough to be able to persistently hold all data that has not yet been uploaded to Amazon S3. When your cache storage has filled up with dirty data (that is, data that has not been uploaded to AWS), application write operations performed to your virtual tape are blocked until more cache storage becomes available. However, application read operations from the virtual tape are still allowed.

Here is some guidance you can follow to help ensure you have adequate storage allocated for your gateway's upload buffer and cache storage:

- **Use Amazon CloudWatch Metrics for the Upload Buffer** – You can proactively avoid the upload buffer from filling up by monitoring the percentage of upload buffer space your gateway-VTL is using over time. Amazon CloudWatch provides metrics such as the `UploadBufferPercentUsed` metric for monitoring your gateway's upload buffer (for more information, see Monitoring the Upload Buffer (p. 355)). You can set an alarm on the `UploadBufferPercentUsed` metric to notify you when upload buffer usage exceeds a threshold. If your upload buffer is getting filled close to capacity, consider adding more upload buffer capacity to the gateway-VTL. For a full list of AWS Storage Gateway metrics, see Understanding Metrics (p. 359).

- **Use Amazon CloudWatch Metrics for Cache Storage** – You can proactively avoid filling up cache storage with dirty data by monitoring how cache storage is being used—particularly, by reviewing cache

misses. Amazon CloudWatch provides usage metrics such as the `CachePercentDirty` and `CacheHitPercent` metrics for monitoring how much of the gateway's cache storage has not been uploaded to Amazon S3. You can set an alarm to trigger a notification to you when the percentage of the cache that's dirty exceeds a threshold or cache hit percentage falls below a threshold. Both can potentially indicate that the cache storage is not adequate for the gateway. Again, for a full list of AWS Storage Gateway metrics, see Understanding Metrics (p. 359).

- **Optimize Your Environment** – If the speed of incoming write operations is too high compared to the outgoing network bandwidth, then the gateway-VTL might never catch up, no matter how much upload buffer capacity you provision. In this case, you might improve the situation by allocating better hardware resources for your gateway. For more information, see Optimizing Gateway-VTL Performance (p. 341).
- **Use the Sizing Formula** – As your application needs change, you should periodically review the recommended formula for sizing the upload buffer. For more information, see Sizing the Upload Buffer (p. 305).

For instructions on working with upload buffer capacity, see Adding and Removing Upload Buffer Capacity (Gateway-VTL) (p. 306).

## Sizing the Upload Buffer

You can determine the size of your upload buffer by using an upload buffer formula. We strongly recommend that you allocate at least 150 GiB of upload buffer. If the formula returns a value less than 150 GiB, use 150 GiB as the amount you allocate to the upload buffer. You can configure up to 2 TiB of upload buffer capacity for each gateway.

> **Note**
> When the upload buffer reaches its capacity, your applications can continue to read from and write data to your storage volumes. However, the gateway-VTL does not write any of your volume data to its upload buffer and does not upload any of this data to AWS until AWS Storage Gateway synchronizes the data stored locally with the copy of the data stored in AWS. This synchronization occurs when the volumes are in BOOTSTRAPPING status.

To estimate the amount of upload buffer, you can determine the expected incoming and outgoing data rates and plug them into the following formula.

- **Rate of Incoming Data—**This rate refers to the application throughput, the rate at which your on-premises applications are writing data to your gateway-VTL over some period of time.
- **Rate of Outgoing Data—**This rate refers to the network throughput, the rate at which your gateway-VTL is able to upload data to AWS. This rate depends on your network speed, utilization, and whether you've enabled bandwidth throttling. This rate should be adjusted for compression. When uploading data to AWS, the gateway-VTL applies data compression where possible. For example, if your application data is text-only, you might get effective compression ratio of about 2:1. However, if you are writing videos, the gateway-VTL might not be able to achieve any data compression, requiring more upload buffer for the gateway.

If your incoming rate is higher than the outgoing rate, or if the formula returns a value less than 150 GiB, we strongly recommend that you allocate at least 150 GiB of upload buffer space.

$$\left(\text{Application Throughput (MB/s)} - \text{Network Throughput to AWS (MB/s)} \times \text{Compression Factor}\right) \times \text{Duration of writes (s)} = \text{Upload Buffer (MB)}$$

For example, assume that your business applications will write text data to your gateway-VTL at a rate of 40 MB per second for 12 hours a day and your network throughput is 12 MB per second. Assuming a compression factor of 2:1 for the text data, you need to allocate approximately 690 GiB of space for the upload buffer.

```
((40 MB/sec) - (12 MB/sec * 2)) * (12 hours * 3600 seconds/hour) = 691200 MB
```

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway-VTL as upload buffer space. Add more upload buffer space as needed using the AWS Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see Monitoring the Upload Buffer (p. 355).

If you decide that you need to change your upload buffer capacity, take one of the following actions.

| To... | Do This... |
|---|---|
| Add more upload buffer capacity to your gateway. | Follow the steps in Adding Upload Buffer Capacity (p. 173). |
| Remove a disk allocated as upload buffer space. | Follow the steps in Removing Upload Buffer Capacity (p. 173). |

### Sizing Cache Storage

The gateway uses the cache storage to provide low-latency access to your recently accessed data. The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer. Generally speaking, you size the cache storage at 1.1 times the upload buffer size. For more information about how to estimate your cache storage size, see Sizing the Upload Buffer (p. 171).

You can initially use this approximation to provision disks for the cache storage. You can then use Amazon CloudWatch operational metrics to monitor the cache storage usage and provision more storage as needed using the console. For information on using the metrics and setting up alarms, see Monitoring Cache Storage (p. 358).

If you decide that you need to increase your gateway's cache storage capacity, follow the steps in Adding and Removing Upload Buffer Capacity (Gateway-VTL) (p. 306).

### Adding and Removing Upload Buffer Capacity (Gateway-VTL)

You can add more upload buffer capacity to your gateway-VTL without interrupting existing gateway-VTL functions. Note that when you add more upload buffer capacity, you can do so with the gateway-VTL VM turned on. However, when you reduce the amount of upload buffer capacity, you must first turn off the VM.

### Adding Upload Buffer or Cache Storage Capacity (Gateway-VTL)

Configuring your gateway-VTL with additional storage as an upload buffer or cache storage is a two-step process.

1. Add a disk of the size you need.
   - For more information about adding disks for upload buffer space on a VMware ESXi host, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23). For a Microsoft Hyper-V host, see Step 2.3: Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 45)
   - For more information about adding disks for cache storage on a VMware ESXi host, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23). For a Microsoft Hyper-V host, see Step 2.3: Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 45).
2. Configure your gateway-VTL to use the additional disk as an upload buffer. For more information, see Adding and Removing Upload Buffer Capacity (Gateway-VTL) (p. 306).

## Removing Upload Buffer Capacity (Gateway-VTL)

As your application needs change and you change the local storage disk configuration for a gateway, you might need to decrease the gateway's upload buffer capacity. Or, a local disk allocated as upload buffer space might fail and you might need to remove that disk from your upload buffer and assign a new local disk. In both cases, you can remove more buffer capacity using the AWS Storage Gateway console.

AWS Storage Gateway does not support removing disks allocated for cache storage.

### To remove a disk allocated as upload buffer

1. On the **Gateway** tab of the AWS Storage Gateway console, click **Configure Local Storage**.

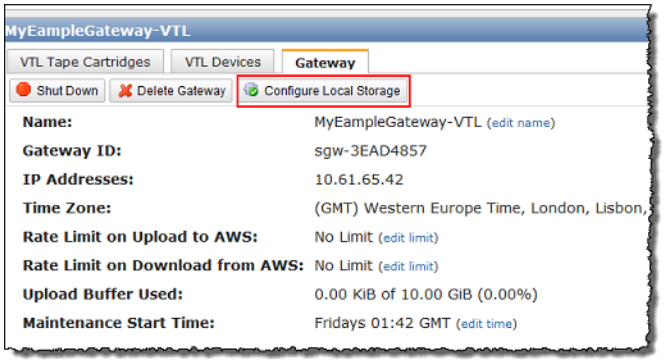


2. In the **Configure Your Activated Gateway** dialog box, note the value of the virtual device node for the local disk to be removed. You can find the node value in the **Local Disks** column. For example, in the following dialog box, the device node **SCSI (0:2)** is highlighted.

   You use the disk's virtual device node in your VM client to ensure that you remove the correct disk.

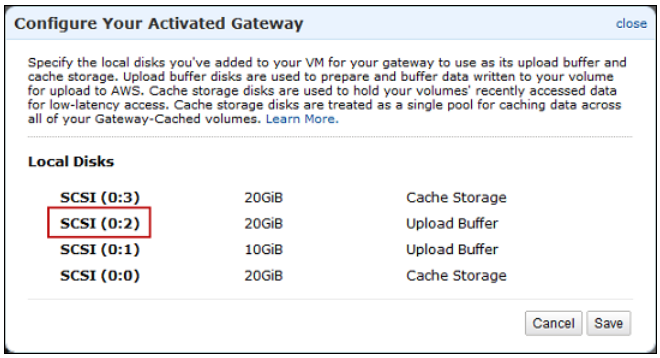
Configure Your Activated Gateway
close
Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage. Upload buffer disks are used to prepare and buffer data written to your volume for upload to AWS. Cache storage disks are used to hold your volumes' recently accessed data for low-latency access. Cache storage disks are treated as a single pool for caching data across all of your Gateway-Cached volumes. Learn More.
Local Disks
SCSI (0:3) 20GiB Cache Storage
SCSI (0:2) 20GiB Upload Buffer
SCSI (0:1) 10GiB Upload Buffer
SCSI (0:0) 20GiB Cache Storage
Cancel
Save

3. Shut down the gateway-VTL by following the steps in the Shutting Down and Starting Your Gateway-VTL (p. 311) procedure.

   **Note**
   Before shutting down the gateway, ensure that no application is writing data to it and that archiving is not in progress. You can check the status of the tapes in the **Status** column of the AWS Storage Gateway console.

4. To remove the underlying local disk, do one of the following procedures.

| For a Gateway Hosted In... | Do This... |
|---|---|
| VMware ESXi | Follow the steps in To remove the local disk allocated as an upload buffer disk (VMware ESXi) (p. 308). |

| For a Gateway Hosted In... | Do This... |
|---|---|
| Microsoft Hyper-V | Follow the steps in To remove the local disk allocated as an upload buffer disk (Microsoft Hyper-V) (p. 309). |

5. In the AWS Storage Gateway console, turn on the gateway.

   **Important**
   After removing a disk used as an upload buffer, you must turn the gateway-VTL back on before adding new disks to the VM.

### To remove the local disk allocated as an upload buffer disk (VMware ESXi)

1. In the vSphere client, right-click the name of your gateway-VTL VM and click **Edit Settings**.
2. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk allocated as upload buffer space, and then click **Remove**.

   Verify that the **Virtual Device Node** value in the **Virtual Machine Properties** dialog box has the same value that you noted previously. Doing this helps ensure that you remove the correct disk.

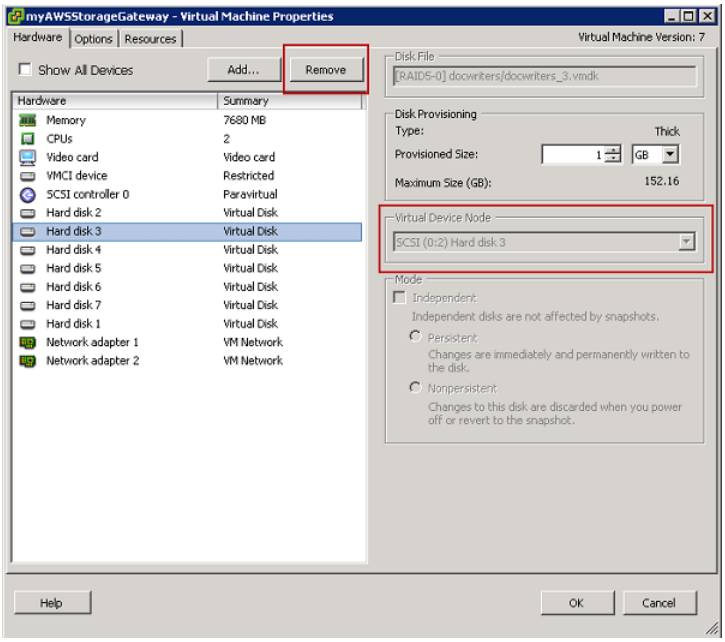


3. Choose an option in the **Removal Options** panel, and then click **OK** to complete the process of removing the disk.

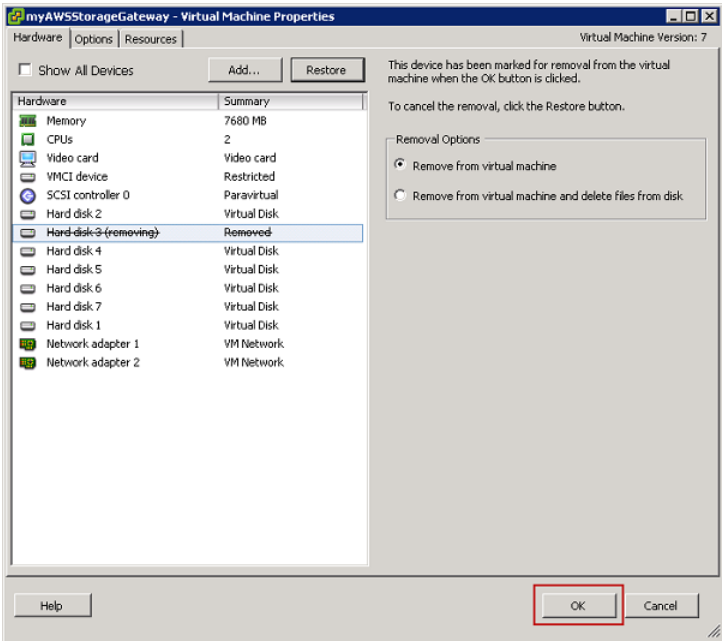


### To remove the local disk allocated as an upload buffer disk (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, right-click the name of your gateway-VTL VM and click **Settings**.
2. In the **Hardware** list of the **Settings** dialog box, select the disk to remove, and then click **Remove**.

   The disks you add to a gateway-VTL are under the **SCSI Controller** entry in the **Hardware** list. Verify that the **Controller** and **Location** value are the same value that you noted previously. Doing this helps ensure that you remove the correct disk.

   The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.

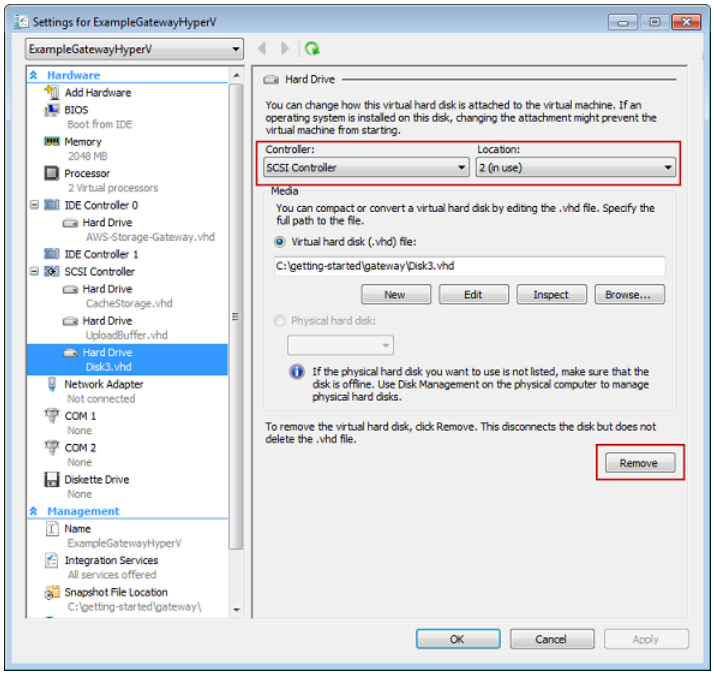

3. Click **OK** to apply the change.

## Configuring an Upload Buffer and Cache Storage (Gateway-VTL)

As part of setting up your gateway-VTL, you allocate storage (that is, disks) for your gateway. You use the suggested sizing formula in deciding the disk sizes. For more information on sizing the upload buffer, see Managing the Upload Buffer and Cache Storage (Gateway-VTL) (p. 304).

For information about the role of the local disks in the overall functioning of the gateway, see Gateway–Virtual Tape Library (Gateway-VTL) Architecture (p. 5). The following diagram highlights the upload buffer and cache storage in the larger picture of gateway-VTL architecture.

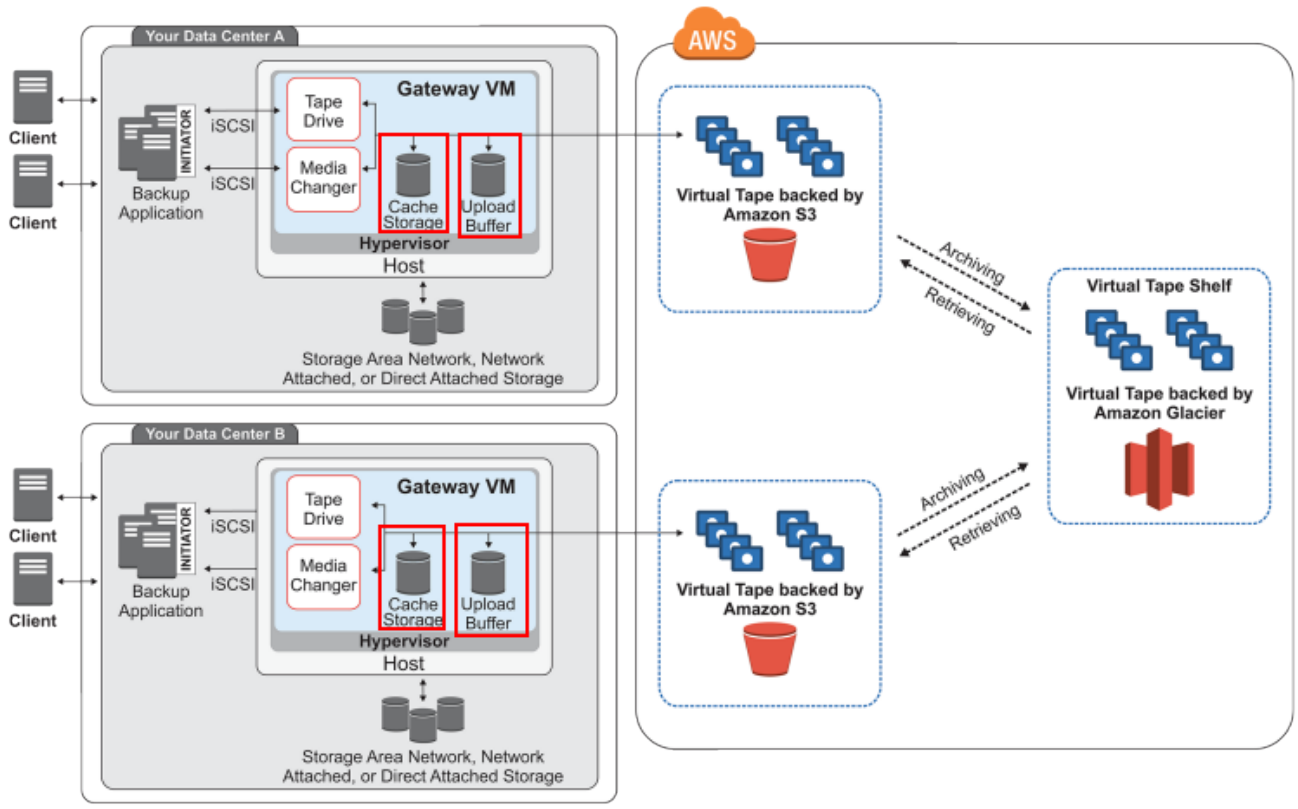


This section shows how you can configure your gateway-VTL to use the allocated disks as an upload buffer and cache storage.

### To configure local disks as an upload buffer and cache storage

1. In the AWS Storage Gateway console, click the **Gateway** tab.
2. Click the **Configure Local Storage** tab.

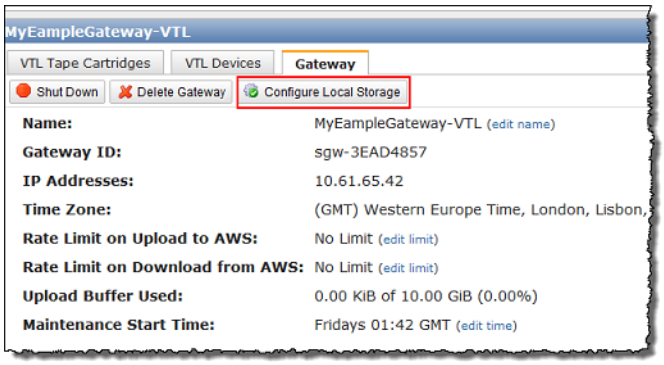


3. In the **Configure Your Activated Gateway** wizard, verify that there are local disks available to configure as an upload buffer and cache storage.

   The wizard shows a list of available disks on your local VM. If there are no local disks available, you must first add disks to your gateway-VTL VM. For more information on doing so, see Step 2.3: Provision Local Disk Storage for the Gateway VM (VMWare) (p. 23) or Step 2.3: Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 45).

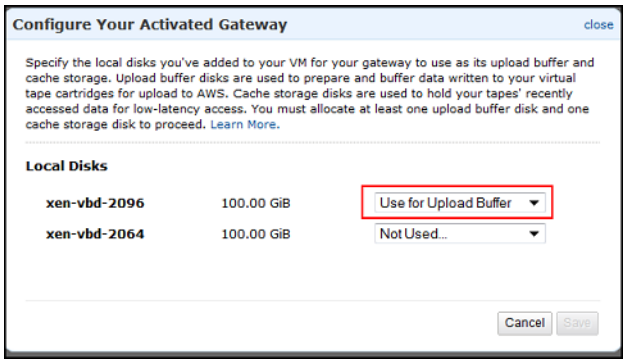
Configure Your Activated Gateway
close
Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage. Upload buffer disks are used to prepare and buffer data written to your virtual tape cartridges for upload to AWS. Cache storage disks are used to hold your tapes' recently accessed data for low-latency access. You must allocate at least one upload buffer disk and one cache storage disk to proceed. Learn More.
Local Disks
xen-vbd-2096
100.00 GiB
Use for Upload Buffer
xen-vbd-2064
100.00 GiB
Not Used...
Cancel
Save

4. For each disk that you want to use, from the drop-down list beside its name, select **Use for Upload Buffer** or **Use for Cache Storage**.

   **Important**
   After configuring these disks, you lose any pre-existing data on the disks.

5. Click **Save** to save your configuration.

## Shutting Down and Starting Your Gateway-VTL

This section discusses shutting down and starting a gateway-VTL. For example, you might need to shut down your gateway VM for host maintenance, such as applying a patch to your hypervisor. Before you power off the gateway VM, you should shut down the gateway.

**Note**
If you shut down your gateway while your backup software is writing or reading from a tape, the write or read task might not succeed. Before shutting down, you should check your backup software and the backup schedule for any tasks in progress.

### To shut down a gateway

1. On the **Navigation** pane of the **AWS Storage Gateway** console, select the gateway.
2. Click **Shut Down**.

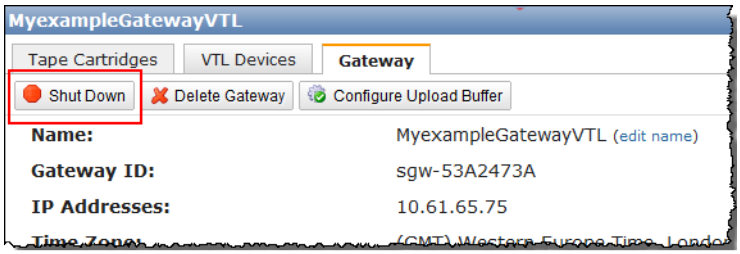


3. In the confirmation dialog box, click **Okay**.

4. In the **Gateway Shut Down** dialog box, click **Close**.

While the gateway is shutting down, you might see a message that your gateway is in the process of shutting down.

5. On the left navigation pane, select the gateway.

A **Restart** button is displayed.

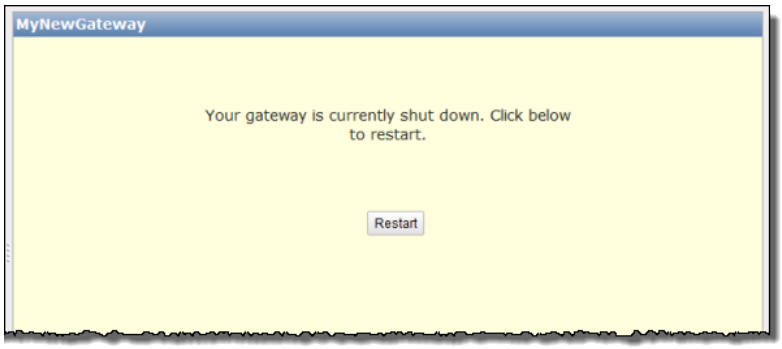


### To start a gateway

1. On the **Navigation** pane of the AWS Storage Gateway console, select the gateway to restart.
2. Click **Restart**.

## Managing Gateway-VTL Updates Using the AWS Storage Gateway Console

AWS Storage Gateway periodically deploys important updates and patches to your gateway that must be applied. AWS will notify you via the AWS Storage Gateway console and via email in advance of any updates to your gateway. Software updates force a restart of your gateway, which typically takes a few minutes to complete. You do not have take any action during this update.

You can choose to let AWS Storage Gateway apply updates according to the maintenance schedule for your gateway, or you can apply the updates yourself. When you deploy and activate your gateway, a default weekly maintenance schedule is set. You can modify this schedule at any time by clicking **edit time** next to **Maintenance Start Time** on the **Gateway** tab of the console. The following example shows the **Gateway** tab with a maintenance message and the **Apply Update Now** button. The date and time the last successful software update was applied to the gateway is displayed in the **Gateway** tab.

> **Important**
> A software update forces your gateway to restart of your gateway. You can minimize the chance of any disruption to your applications by increasing your iSCSI initiators' timeouts. For more information about increasing iSCSI initiator timeouts for Windows and Linux, see Customizing Your Windows iSCSI Settings (p. 112) and Customizing Your Linux iSCSI Settings (p. 115).

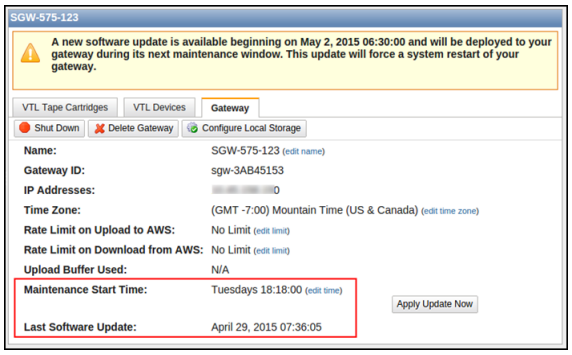
SGW-575-123
A new software update is available beginning on May 2, 2015 06:30:00 and will be deployed to your gateway during its next maintenance window. This update will force a system restart of your gateway.
VTL Tape Cartridges
VTL Devices
Gateway
Shut Down
Delete Gateway
Configure Local Storage
Name:
SGW-575-123 (edit name)
Gateway ID:
sgw-3AB45153
IP Addresses:
Time Zone:
(GMT -7:00) Mountain Time (US & Canada) (edit time zone)
Rate Limit on Upload to AWS:
No Limit (edit limit)
Rate Limit on Download from AWS:
No Limit (edit limit)
Upload Buffer Used:
N/A
Maintenance Start Time:
Tuesdays 18:18:00 (edit time)
Apply Update Now
Last Software Update:
April 29, 2015 07:36:05

## Updating Bandwidth Rate Limits for Your Gateway-VTL

You can limit (or throttle) the upload throughput from the gateway to AWS or the download throughput from your AWS to your gateway. Using bandwidth throttling helps you to control the amount of network bandwidth used by your gateway. By default, an activated gateway has no rate limits on upload or download.

You can specify the rate limit by using the AWS Management Console, or programmatically by using either the AWS Storage Gateway API (see UpdateBandwidthRateLimit) or an AWS Software Development Kit (SDK). The ability to change throttling programmatically lets you change limits automatically throughout the day—for example, by scheduling tasks to change the bandwidth. This section describes how you can change these limits by using the AWS Storage Gateway console. For information about changing bandwidth rate limits programmatically, see the following topics.

**Topics**

- Updating Gateway-VTL Bandwidth Rate Limits Using the AWS SDK for Java (p. 314)
- Updating Gateway-VTL Bandwidth Rate Limits Using the AWS SDK for .NET (p. 316)
- Updating Gateway-VTL Bandwidth Rate Limits Using the AWS Tools for Windows PowerShell (p. 318)

### To change a gateway's bandwidth throttling using the console

1. On the **Navigation** pane of the AWS Storage Gateway console, select the gateway you want to manage.
2. On the right pane, click the **Gateway** tab.
3. Next to the limit you want to change, click **edit limit**.

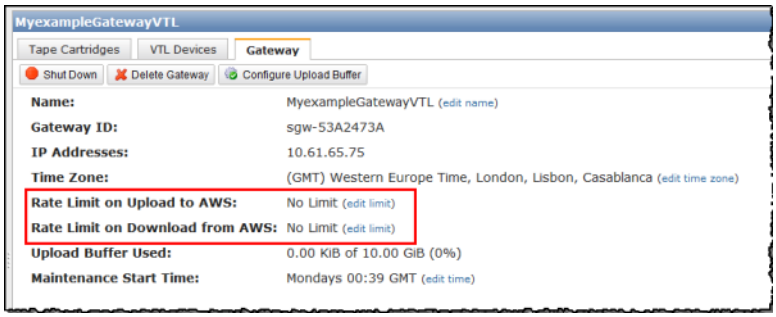


4. In the **Edit Rate Limits** dialog box, enter new limit values, and then click **Save**.

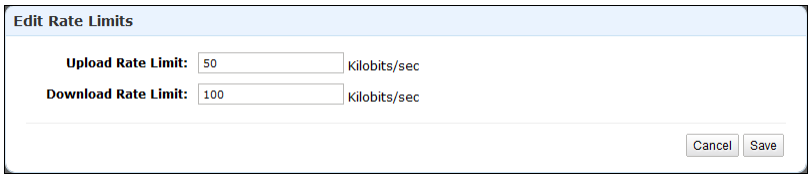


## Updating Gateway-VTL Bandwidth Rate Limits Using the AWS SDK for Java

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth rate limits using the AWS SDK for Java. To use the example code, you should be familiar with running a Java console application. For more information, see Getting Started in the *AWS SDK for Java Developer Guide*.

### Example : Updating Gateway Bandwidth Limits Using the AWS SDK for Java

The following Java code example updates a gateway's bandwidth rate limits. You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints you can use with AWS Storage Gateway, see Regions and Endpoints in the *AWS General Reference*.

```
import java.io.IOException;

import com.amazonaws.AmazonClientException;
import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.storagegateway.AWSStorageGatewayClient;
import com.amazonaws.services.storagegateway.model.UpdateBandwidthRateLimitRe
quest;
import com.amazonaws.services.storagegateway.model.UpdateBandwidthRateLimitRes
ult;


public class UpdateBandwidthExample {

    public static AWSStorageGatewayClient sgClient;

    // The gatewayARN
    public static String gatewayARN = "*** provide gateway ARN ***";

    // The endpoint
   static String serviceURL = "https://storagegateway.us-east-1.amazonaws.com";


    // Rates
    static long uploadRate = 51200;  // Bits per second, minimum 51200
    static long downloadRate = 102400;   // Bits per second, minimum 102400

    public static void main(String[] args) throws IOException {

        // Create a storage gateway client
        sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(
                ListDeleteVolumeSnapshotsExample.class.getResourceAs
Stream("AwsCredentials.properties")));
        sgClient.setEndpoint(serviceURL);

        UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

    }

    private static void UpdateBandwidth(String gatewayARN2, long uploadRate2,
            long downloadRate2) {
        try
        {
            UpdateBandwidthRateLimitRequest updateBandwidthRateLimitRequest =
                new UpdateBandwidthRateLimitRequest()
                .withGatewayARN(gatewayARN)
                .withAverageDownloadRateLimitInBitsPerSec(downloadRate)
                .withAverageUploadRateLimitInBitsPerSec(uploadRate);

            UpdateBandwidthRateLimitResult updateBandwidthRateLimitResult =
sgClient.updateBandwidthRateLimit(updateBandwidthRateLimitRequest);
           String returnGatewayARN = updateBandwidthRateLimitResult.getGatewa
```

```
yARN();
            System.out.println("Updated the bandwidth rate limits of " + re
turnGatewayARN);
            System.out.println("Upload bandwidth limit = " + uploadRate + "
bits per second");
            System.out.println("Download bandwidth limit = " + downloadRate +
" bits per second");
        }
        catch (AmazonClientException ex)
        {
            System.err.println("Error updating gateway bandwith.\n" + ex.to
String());
        }
    }
}
```

### Updating Gateway-VTL Bandwidth Rate Limits Using the AWS SDK for .NET

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth rate limits using the AWS Software Development Kit (SDK) for .NET. To use the example code, you should be familiar with running a .NET console application. For more information, see Getting Started in the *AWS SDK for .NET Developer Guide*.

### Example : Updating Gateway Bandwidth Limits Using the AWS SDK for .NET

The following C# code example updates a gateway's bandwidth rate limits. You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints you can use with AWS Storage Gateway, see Regions and Endpoints in the *AWS General Reference*.

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using Amazon.StorageGateway;
using Amazon.StorageGateway.Model;

namespace AWSStorageGateway
{
    class UpdateBandwidthExample
    {
        static AmazonStorageGatewayClient sgClient;
        static AmazonStorageGatewayConfig sgConfig;

        // The gatewayARN
        public static String gatewayARN = "*** provide gateway ARN ***";

        // The endpoint
        static String serviceURL = "https://storagegateway.us-east-1.amazon
aws.com";

        // Rates
        static long uploadRate = 51200;  // Bits per second, minimum 51200
       static long downloadRate = 102400;   // Bits per second, minimum 102400


        public static void Main(string[] args)
        {
            // Create a storage gateway client
            sgConfig = new AmazonStorageGatewayConfig();
            sgConfig.ServiceURL = serviceURL;
            sgClient = new AmazonStorageGatewayClient(sgConfig);

            UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

            Console.WriteLine("\nTo continue, press Enter.");
            Console.Read();
        }

        public static void UpdateBandwidth(string gatewayARN, long uploadRate,
 long downloadRate)
        {
            try
            {
               UpdateBandwidthRateLimitRequest updateBandwidthRateLimitRequest
 =
                    new UpdateBandwidthRateLimitRequest()
                    .WithGatewayARN(gatewayARN)
                    .WithAverageDownloadRateLimitInBitsPerSec(downloadRate)
                    .WithAverageUploadRateLimitInBitsPerSec(uploadRate);
```

```
             UpdateBandwidthRateLimitResponse updateBandwidthRateLimitResponse
 = sgClient.UpdateBandwidthRateLimit(updateBandwidthRateLimitRequest);
                String returnGatewayARN = updateBandwidthRateLimitResponse.Up
dateBandwidthRateLimitResult.GatewayARN;
                Console.WriteLine("Updated the bandwidth rate limits of " +
returnGatewayARN);
                Console.WriteLine("Upload bandwidth limit = " + uploadRate + "
 bits per second");
                Console.WriteLine("Download bandwidth limit = " + downloadRate
 + " bits per second");
            }
            catch (AmazonStorageGatewayException ex)
            {
                Console.WriteLine("Error updating gateway bandwith.\n" +
ex.ToString());
            }
        }
    }
}
```

## Updating Gateway-VTL Bandwidth Rate Limits Using the AWS Tools for Windows PowerShell

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth rate limits using the AWS Tools for Windows PowerShell. To use the example code, you should be familiar with running a PowerShell script. For more information, see Getting Started in the *AWS Tools for Windows PowerShell User Guide*.

### Example : Updating Gateway Bandwidth Limits Using the AWS Tools for Windows PowerShell

The following PowerShell script example updates a gateway's bandwidth rate limits. You need to update the script and provide your gateway Amazon Resource Name (ARN), and the upload and download limits.

```
<#
.DESCRIPTION
    Update Gateway bandwidth limits.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from http://console.amazonaws.cn/powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
   For more info see, http://docs.amazonaws.cn/powershell/latest/userguide/spe
cifying-your-aws-credentials.html

.EXAMPLE
    powershell.exe .\SG_UpdateBandwidth.ps1
#>

$UploadBandwidthRate = 51200
$DownloadBandwidthRate = 102400
$gatewayARN = "*** provide gateway ARN ***"

#Update Bandwidth Rate Limits
Update-SGBandwidthRateLimit -GatewayARN $gatewayARN `
                            -AverageUploadRateLimitInBitsPerSec $UploadBandwid
thRate `
                           -AverageDownloadRateLimitInBitsPerSec $DownloadBand
widthRate

$limits =  Get-SGBandwidthRateLimit -GatewayARN $gatewayARN

Write-Output("`nGateway: " + $gatewayARN);
Write-Output("`nNew Upload Rate: " + $limits.AverageUploadRateLimitInBitsPerSec)
Write-Output("`nNew Download Rate: " + $limits.AverageDownloadRateLimitInBitsPer
Sec)
```

## Disabling Your Gateway-VTL

Disabling a gateway locks down the virtual tapes. That is, any data that you might write to these tapes after disabling the gateway will not be sent to AWS. If you have a failed gateway and you want to recover the tapes to another gateway, you must first disable the failed gateway. For more information about recovering tapes, see Recovering a Virtual Tape from a Malfunctioning Gateway-VTL (p. 282).

Note that you can disable a gateway by using the console only when the gateway is not connected to AWS. If the gateway is connected to AWS, you will not see the **Disable Gateway** button in the console.

### To disable your gateway

1. In the AWS Storage Gateway console, select the failed gateway, and then click **Disable Gateway**.

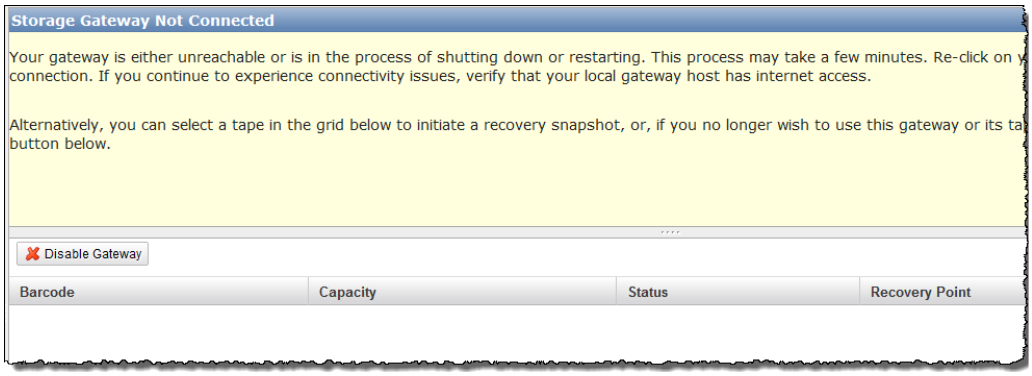
Storage Gateway Not Connected
Your gateway is either unreachable or is in the process of shutting down or restarting. This process may take a few minutes. Re-click on y
connection. If you continue to experience connectivity issues, verify that your local gateway host has internet access.
Alternatively, you can select a tape in the grid below to initiate a recovery snapshot, or, if you no longer wish to use this gateway or its ta
button below.
Disable Gateway
Barcode
Capacity
Status
Recovery Point

2. In the dialog box that appears, click **OK**.

## Deleting Your Gateway-VTL

Deleting deactivates your gateway-VTL. You can delete a gateway-VTL either in the AWS Storage Gateway console or programmatically. When you delete a gateway, you can encounter two scenarios.

- **The gateway-VTL is connected to AWS –** If the gateway-VTL is connected to AWS and you delete the gateway, the iSCSI targets associated with the gateway (that is, the virtual tape drives and media changer) will no longer be available.
- **The gateway-VTL is not connected to AWS –** If the gateway-VTL is not connected to AWS, for example if the underlying VM is turned off or your network is down, then you cannot delete the gateway. If you attempt to do so, after your environment is back up and running, you might have a gateway-VTL running on-premises with available iSCSI targets. However, no gateway-VTL data will be uploaded to or downloaded from AWS.

> **Note**
> If the gateway-VTL you want to delete is not functioning, you must first disable the gateway-VTL and delete tapes that are in RETRIEVING or RETRIEVED status before you delete the gateway-VTL.

After you delete the gateway-VTL in the console, you should remove the associated VM from your host.

After you delete the gateway, you are no longer charged for it. If you have tapes archived in VTS, those tapes remain and you continue to pay for the storage. Before you delete a gateway, you must remove tapes from the VTL as follows.

- If the tape status is RETRIEVING, you must cancel the retrieval process. For instructions, see Canceling Tape Retrieval (p. 281).
- If the tape status is RETRIEVED, eject the tape using your backup software. For instructions, see Archiving the Tape (p. 259).

Before deleting a gateway-VTL, you must cancel all tape retrieval operations and eject all retrieved tapes.

### To delete a gateway-VTL

1. On the navigation pane of the **AWS Storage Gateway** console, select the gateway-VTL you want to delete.
2. In the **Gateway** tab, click **Delete Gateway**.

> **Important**
> Be sure that there are no applications currently writing to the gateway-VTL. If you delete the gateway-VTL while it is in use, data loss can occur.

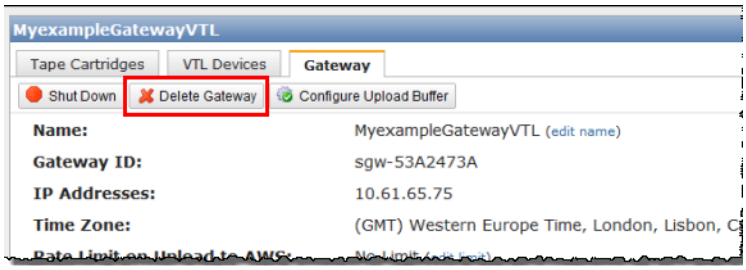


3. Click **OK** to confirm the deletion.

   At this point, the deleted gateway-VTL is no longer activated. However, the gateway-VTL VM still exists in your virtualization environment. To remove the VM, use the VMware vSphere client or Microsoft Hyper-V Manager to connect to the host and remove the VM.

# Maintenance Tasks When Running a Gateway-VTL On-Premises

For a gateway-VTL deployed on-premises, this section describes the maintenance tasks you can perform using the local gateway console.

**Topics**


- Logging In to Your Gateway-VTL's Local Console (p. 321)
- Routing Your Gateway-VTL Through a Proxy (p. 325)
- Configuring Your Gateway-VTL to Use a Static IP Address (p. 326)
- Testing Your Gateway-VTL's Connectivity to the Internet (p. 329)
- Synchronizing Your Gateway-VTL's VM Time (p. 330)
- Running Storage Gateway Commands on the Local Console (p. 332)
- Updating Your System (p. 333)
- Viewing Your Gateway-VTL System Resource Errors (p. 334)
- Configuring Your Gateway-VTL for Multiple NICs (p. 335)
- Optimizing Gateway-VTL Performance (p. 341)


## Logging In to Your Gateway-VTL's Local Console

Some gateway-VTL maintenance tasks require that you log in to your gateway's local console. You can access the local console through your hypervisor client software. The user is *sguser* and the password is *sgpassword*. These login credentials give you access to menus where you can configure gateway network settings. After you log in, the local console looks like the following screenshot.

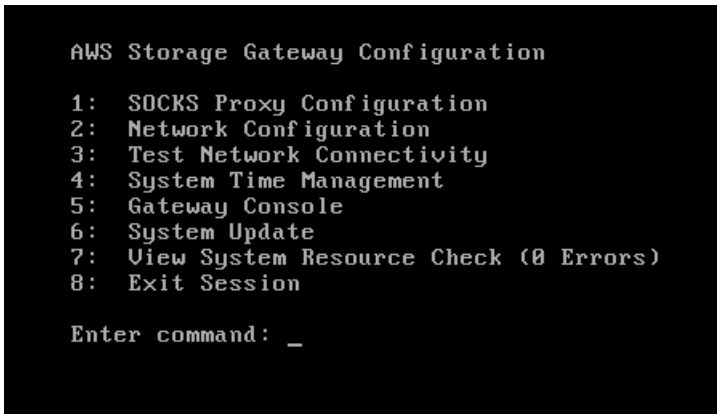


In this topic, we show you how to access the local console of a gateway-VTL hosted in these hypervisors:

- VMware ESXi—for more information, see To access your gateway-VTL's local console (VMware ESXi) (p. 322).
- Microsoft Hyper-V—for more information, see To access your gateway-VTL's local console (Microsoft Hyper-V) (p. 323).

### To access your gateway-VTL's local console (VMware ESXi)

1. In the VMware vSphere client, select your gateway-VTL VM.
2. Ensure that the gateway-VTL is turned on.

   **Note**
   If your gateway-VTL VM is turned on, a green arrow icon appears with the VM icon, as shown in the following example. If your gateway-VTL VM is not turned on, you can turn it on by clicking the green **Power On** icon on the **Toolbar** menu.

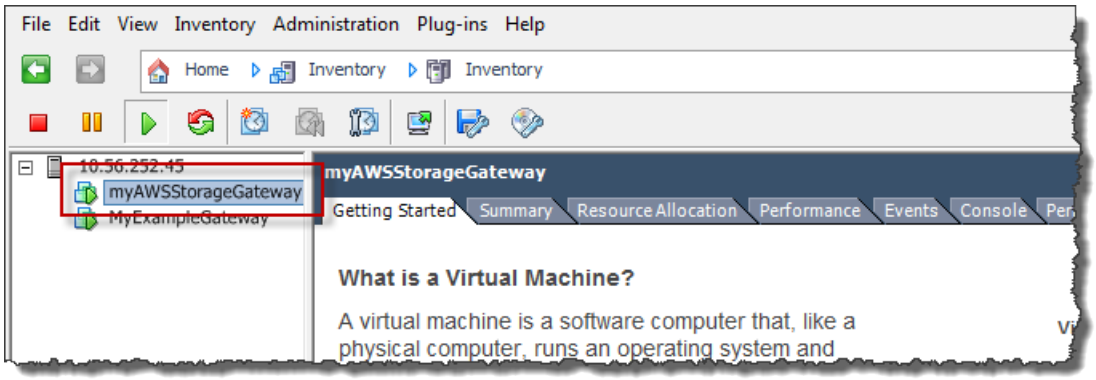
File Edit View Inventory Administration Plug-ins Help
Home Inventory Inventory
10.56.252.45
myAWSStorageGateway
MyExampleGateway
myAWSStorageGateway
Getting Started Summary Resource Allocation Performance Events Console
What is a Virtual Machine?
A virtual machine is a software computer that, like a physical computer, runs an operating system and

3. Click the **Console** tab.

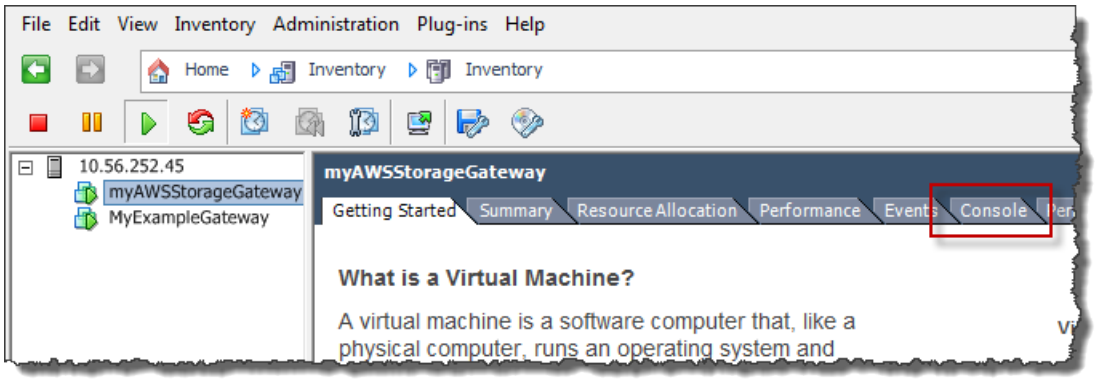
File Edit View Inventory Administration Plug-ins Help
Home Inventory Inventory
10.56.252.45
myAWSStorageGateway
MyExampleGateway
myAWSStorageGateway
Getting Started Summary Resource Allocation Performance Events Console
What is a Virtual Machine?
A virtual machine is a software computer that, like a physical computer, runs an operating system and

4. After a few moments, the VM is ready for you to log in.

   **Note**
   To release the cursor from the console window, press **Ctrl+Alt**.

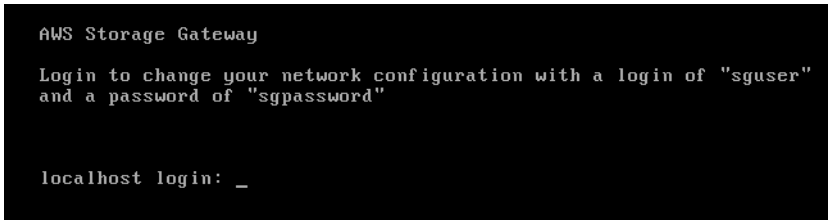
AWS Storage Gateway
Login to change your network configuration with a login of "sguser" and a password of "sgpassword"
localhost login: _

5. To log in, continue to the procedure To log in to the gateway-VTL's local console (p. 324).

### To access your gateway-VTL's local console (Microsoft Hyper-V)

1. In the **Virtual Machines** list of the Microsoft Hyper-V Manager, select your gateway-VTL VM.
2. Ensure the gateway-VTL is turned on.

   **Note**
   If your gateway-VTL VM is turned on, `Running` is displayed as the **State** of the VM, as shown in the following screenshot. If your gateway-VTL VM is not turned on, you can turn it on by clicking **Start** in the **Actions** pane.

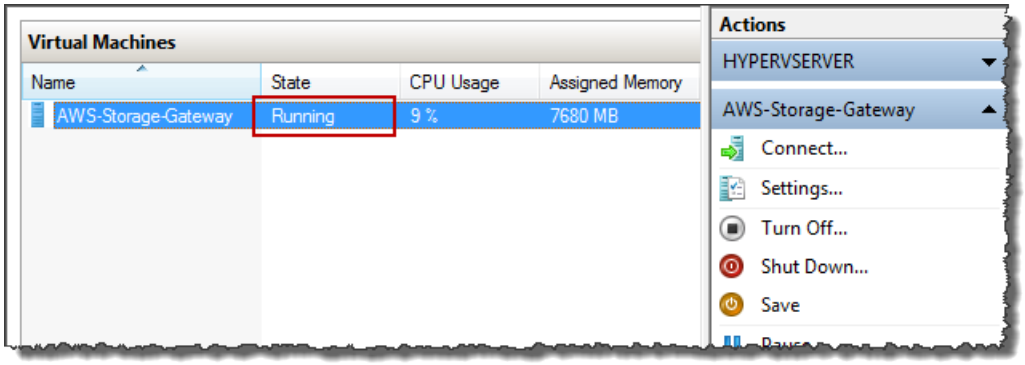


3. In the **Actions** pane, select **Connect**.

   The **Virtual Machine Connection** window appears. If an authentication window appears, type the user name and password provided to you by the hypervisor administrator.

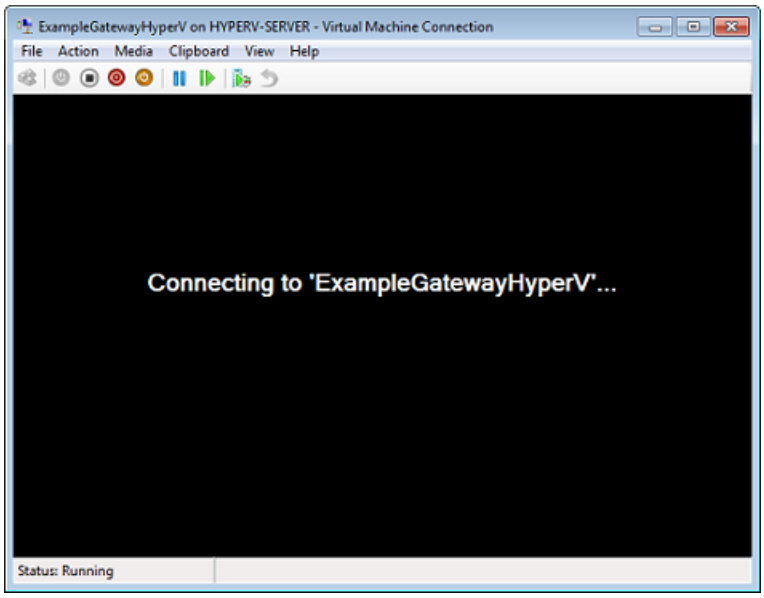
ExampleGatewayHyperV on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
Connecting to 'ExampleGatewayHyperV'...
Status: Running

4. After a few moments, the VM is ready for you to log in.

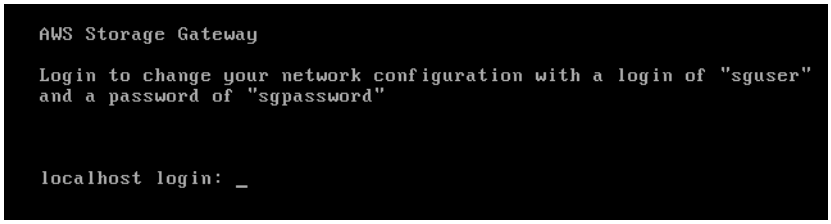
AWS Storage Gateway
Login to change your network configuration with a login of "sguser"
and a password of "sgpassword"
localhost login: _

5. To log in, continue to the procedure To log in to the gateway-VTL's local console (p. 324).

### To log in to the gateway-VTL's local console

- In the login screen, log in to the VM with the user name and password specified in your hypervisor's console window.

After you log in, you will see the **AWS Storage Gateway Configuration** main menu, as shown in the following screenshot.

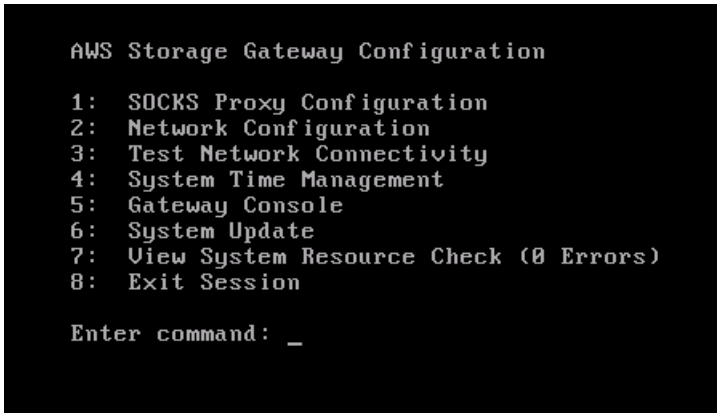

| To... | See... |
|---|---|
| Configure a SOCKS proxy for your gateway-VTL | Routing Your Gateway-VTL Through a Proxy (p. 325). |
| Configure static IP addresses for your gateway's interfaces | Configuring Your Gateway-VTL to Use a Static IP Address (p. 326). |
| Test network connectivity | Testing Your Gateway-VTL's Connectivity to the Internet (p. 329). |
| Manage VM time | Synchronizing Your Gateway-VTL's VM Time (p. 330). |
| Run Storage Gateway console commands | Running Storage Gateway Commands on the Local Console (p. 332). |
| Update your system | Updating Your System (p. 333). |

To exit the configuration session, type **7** to exit the menu.

## Routing Your Gateway-VTL Through a Proxy

AWS Storage Gateway supports the configuration of a Socket Secure Version 5 (SOCKS5) proxy between your gateway-VTL and AWS.

> **Note**
> The only proxy configuration AWS Storage Gateway supports is SOCKS5.

If your gateway-VTL must use a proxy server to communicate to the Internet, then you need to configure the SOCKS proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy. After you do this, AWS Storage Gateway will route all HyperText Transfer Protocol Secure (HTTPS) traffic through your proxy server.

### To route your gateway-VTL Internet traffic through a local proxy server

1. Log in to your gateway's local console. For instructions, see Logging In to Your Gateway-VTL's Local Console (p. 321).
2. On the **AWS Storage Gateway Configuration** main menu, type **1** to begin configuring the SOCKS proxy.

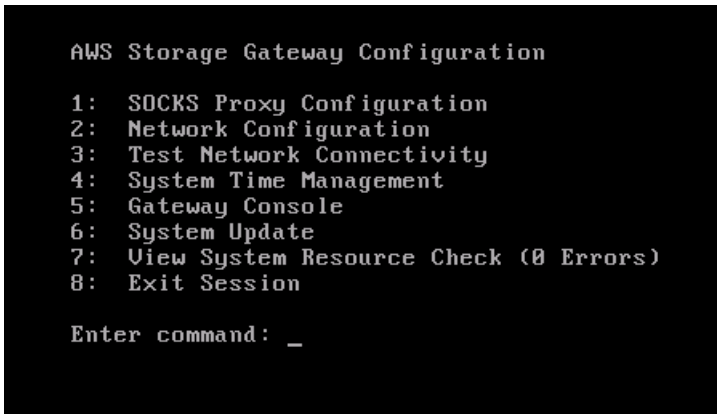


3. Choose one of the following options in the **AWS Storage Gateway SOCKS Proxy Configuration** menu:

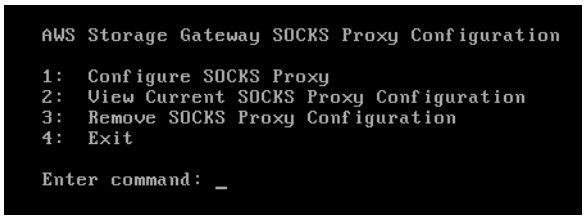
AWS Storage Gateway SOCKS Proxy Configuration
1: Configure SOCKS Proxy
2: View Current SOCKS Proxy Configuration
3: Remove SOCKS Proxy Configuration
4: Exit
Enter command: _

| To... | Do this... |
|---|---|
| Configure a SOCKS proxy | Type option **1**.<br><br>You will need to supply a host name and port to complete configuration. |
| View the current SOCKS proxy configuration | Type option **2**.<br><br>If a SOCKS proxy is not configured, the message `SOCKS Proxy not configured` is displayed. If a SOCKS proxy is configured, the host name and port of the proxy are displayed. |
| Remove a SOCKS proxy configuration | Type option **3**.<br><br>The message `SOCKS Proxy Configuration Removed` is displayed. |
| Exit this menu and return to the previous menu | Type option **4**. |

## Configuring Your Gateway-VTL to Use a Static IP Address

The default network configuration for the gateway-VTL is Dynamic Host Configuration Protocol (DHCP). With DHCP, your gateway-VTL is automatically assigned an IP address. In some cases, you might need to manually assign your gateway's IP as a static IP address. This topic explains how to do so.

### To configure your gateway-VTL to use static IP addresses

1. Log in to your gateway's local console. For instructions, see Logging In to Your Gateway-VTL's Local Console (p. 321).
2. In the **AWS Storage Gateway Configuration** main menu, type option **2** to begin configuring a static IP address.

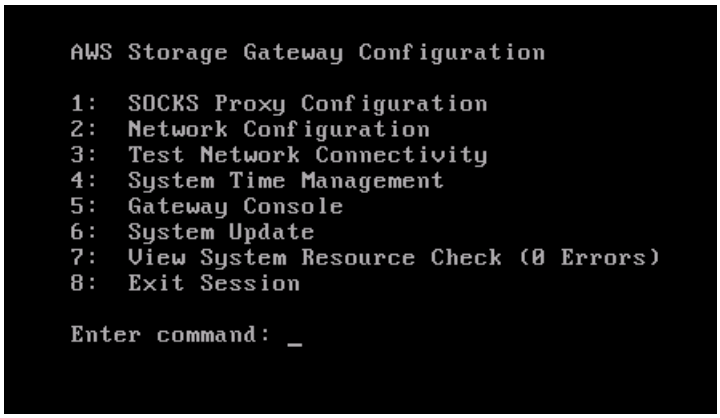


3. Choose one of the following options in the **AWS Storage Gateway Static IP Address Configuration** menu:

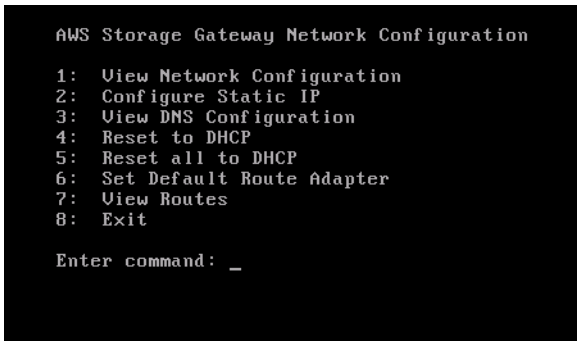


| To... | Do this... |
|---|---|
| View your gateway's network configuration | Type option **1**.<br><br>A list of adapter names appears, and you are prompted to enter an adapter name—for example, **`eth0`**. If the adapter you specify is in use, the following information about the adapter is displayed:<br><br>• Media access control (MAC) address<br>• IP address<br>• Netmask<br>• Gateway IP address<br>• DHCP enabled status<br><br>You use the same adapter name when you configure a static IP address (option 2) or set your gateway VTL's default route adapter (option 5). |

| To... | Do this... |
| --- | --- |
| Configure a static IP address for your gateway | Type option **2**.<br><br>You are prompted to type the following information to configure a static IP:<br><br>• Network adapter name<br>• IP address<br>• Netmask<br>• Default gateway-VTL address<br>• Primary Domain Name Service (DNS) address<br>• Secondary DNS address<br><br>**Important**<br>If your gateway-VTL has already been activated, you must shut it down and restart it from the AWS Storage Gateway console for the settings to take effect. For more information, see Shutting Down and Starting Your Gateway-VTL (p. 311).<br><br>If your gateway-VTL uses more than one network interface, you must set all enabled interfaces to use DHCP or static IP addresses.<br><br>For example, suppose your gateway-VTL VM uses two interfaces configured as DHCP. If you later set one interface to a static IP, the other interface is disabled. To enable the interface in this case, you must set it to a static IP.<br><br>If both interfaces are initially set to use static IP addresses and you then set the gateway-VTL to use DHCP, both interfaces will use DHCP. |
| View your gateway's DNS configuration | Type option **3**.<br><br>The IP addresses of the primary and secondary DNS name servers are displayed. |
| Reset your gateway's network configuration to DHCP | Type option **4**.<br><br>All network interfaces are set to use DHCP.<br><br>**Important**<br>If your gateway-VTL has already been activated, you must shut it down and restart it from the AWS Storage Gateway console for the settings to take effect. For more information, see Shutting Down and Starting Your Gateway-VTL (p. 311). |

| To... | Do this... |
| --- | --- |
| Set your gateway's default route adapter | Type option **5**.<br><br>The available adapters for your gateway-VTL are shown, and you are prompted to select one of the adapters—for example, **`eth0`**. |
| View routing tables | Type option **6**.<br><br>The default route of your gateway-VTL is displayed. |
| Exit this menu and return to the previous menu | Type option **7**. |

## Testing Your Gateway-VTL's Connectivity to the Internet

You can use your gateway's local console to test your Internet connection. This test can be useful when you are troubleshooting network issues with your gateway-VTL. The network connectivity test does not take into account any SOCKS proxy you might have configured.

### To test your gateway-VTL's connection to the Internet

1. Log in to your gateway-VTL's local console. For instructions, see Logging In to Your Gateway-VTL's Local Console (p. 321).
2. In the **AWS Storage Gateway Configuration** main menu, type option **3** to begin testing network connectivity.

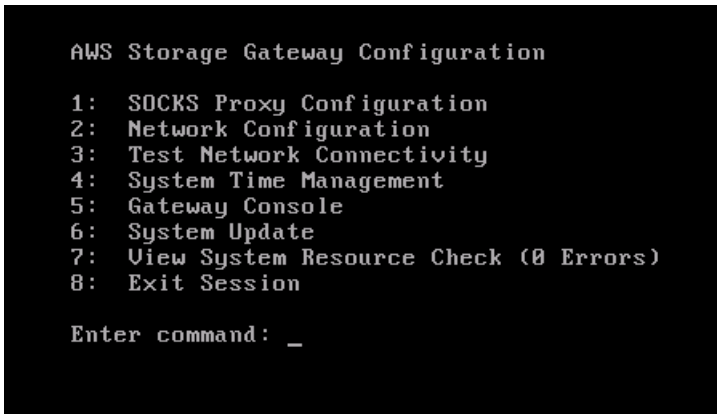


   The console displays the available regions.
3. Select the region you want to test, as shown in the following screenshot.

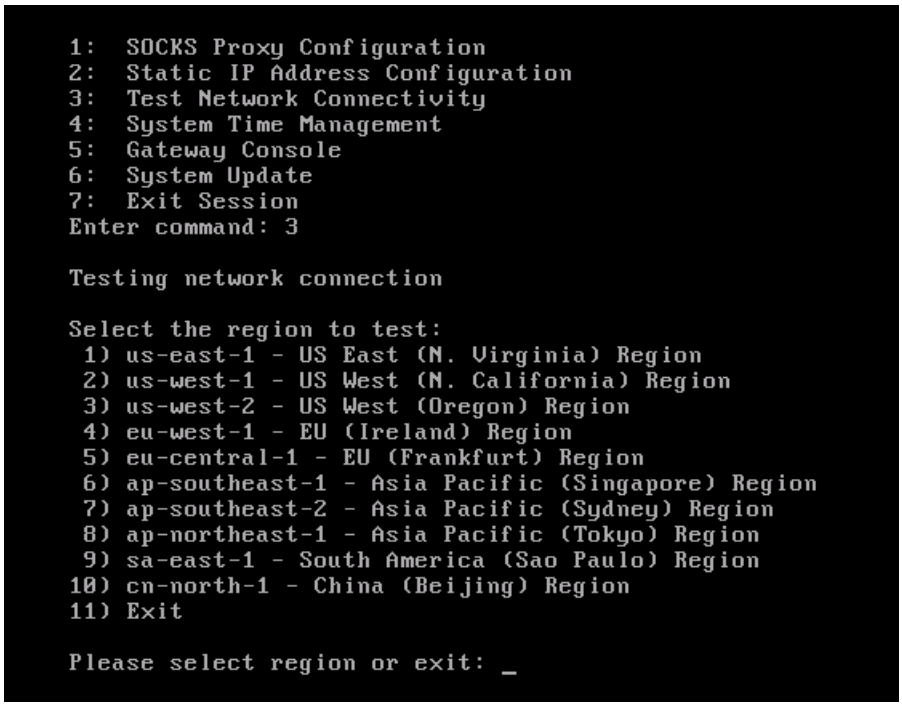
1: SOCKS Proxy Configuration
2: Static IP Address Configuration
3: Test Network Connectivity
4: System Time Management
5: Gateway Console
6: System Update
7: Exit Session
Enter command: 3
Testing network connection
Select the region to test:
1) us-east-1 - US East (N. Virginia) Region
2) us-west-1 - US West (N. California) Region
3) us-west-2 - US West (Oregon) Region
4) eu-west-1 - EU (Ireland) Region
5) eu-central-1 - EU (Frankfurt) Region
6) ap-southeast-1 - Asia Pacific (Singapore) Region
7) ap-southeast-2 - Asia Pacific (Sydney) Region
8) ap-northeast-1 - Asia Pacific (Tokyo) Region
9) sa-east-1 - South America (Sao Paulo) Region
10) cn-north-1 - China (Beijing) Region
11) Exit
Please select region or exit: _

Each endpoint in the selected region displays either a PASSED or FAILED message, as shown following.

| Message | Description |
|---|---|
| [ PASSED ] | AWS Storage Gateway has Internet connectivity. |
| [ FAILED ] | AWS Storage Gateway does not have Internet connectivity. |

## Synchronizing Your Gateway-VTL's VM Time

After your gateway-VTL is deployed and running, in some scenarios the gateway-VTL VM's time can drift. For example, if there is a prolonged network outage and your hypervisor host and gateway-VTL do not get time updates, then the gateway-VTL VM's time will be different from the true time. When there is a time drift, a discrepancy occurs between the stated times when operations such as gateway-VTL updates or snapshots occur and the actual times that the operations occur.

For a gateway-VTL deployed on VMware ESXi, setting the hypervisor host time and synchronizing the VM time to the host is sufficient to avoid time drift. For more information, see Step 2.2.3: Synchronize VM Time with Host Time (p. 20). For a gateway-VTL deployed on Microsoft Hyper-V, you should periodically check your VM's time using the procedure described following.

### To view and synchronize the time of a Hyper-V gateway-VTL VM to an NTP server

1. Log in to your gateway's local console. For instructions, see Logging In to Your Gateway-VTL's Local Console (p. 321).
2. On the **AWS Storage Gateway Configuration** main menu, type option **4** for **System Time Management**.

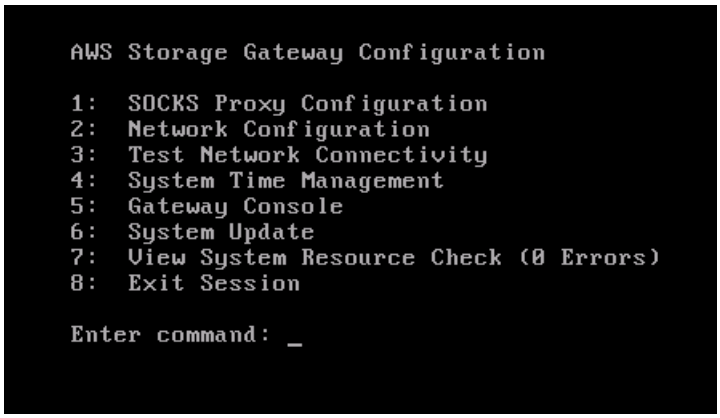


3. On the **System Time Management** menu, type option **1** for **View and Synchronize System Time**.

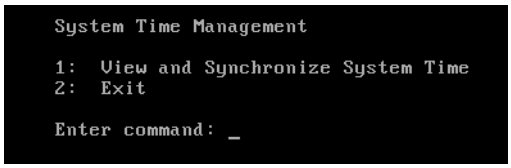


4. If the result indicates that you should synchronize your VM's time to the Network Time Protocol (NTP) time, type **y**. Otherwise, enter **n**. We recommend synchronizing your VM's time to the NTP time if the time differs by more than 60 seconds.

   If you type **y** to synchronize, the synchronization might take a few moments.

   The following screenshot shows a VM that does not require time synchronization.

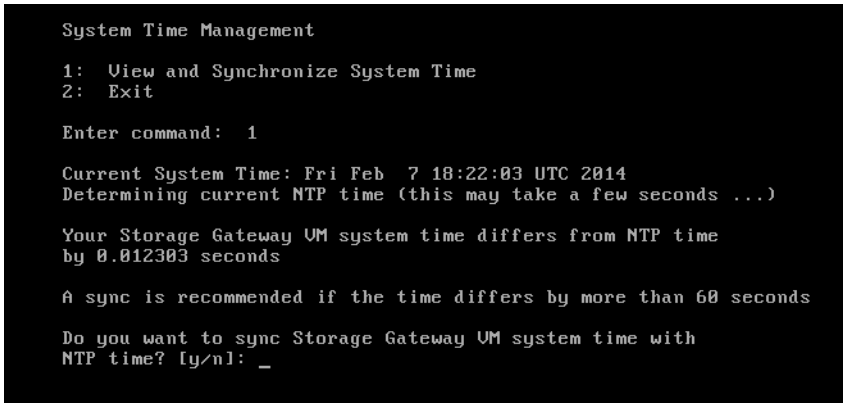
System Time Management
1: View and Synchronize System Time
2: Exit
Enter command: 1
Current System Time: Fri Feb 7 18:22:03 UTC 2014
Determining current NTP time (this may take a few seconds ...)
Your Storage Gateway VM system time differs from NTP time
by 0.012303 seconds
A sync is recommended if the time differs by more than 60 seconds
Do you want to sync Storage Gateway VM system time with
NTP time? [y/n]: _

   The following screenshot shows a VM that does require time synchronization.

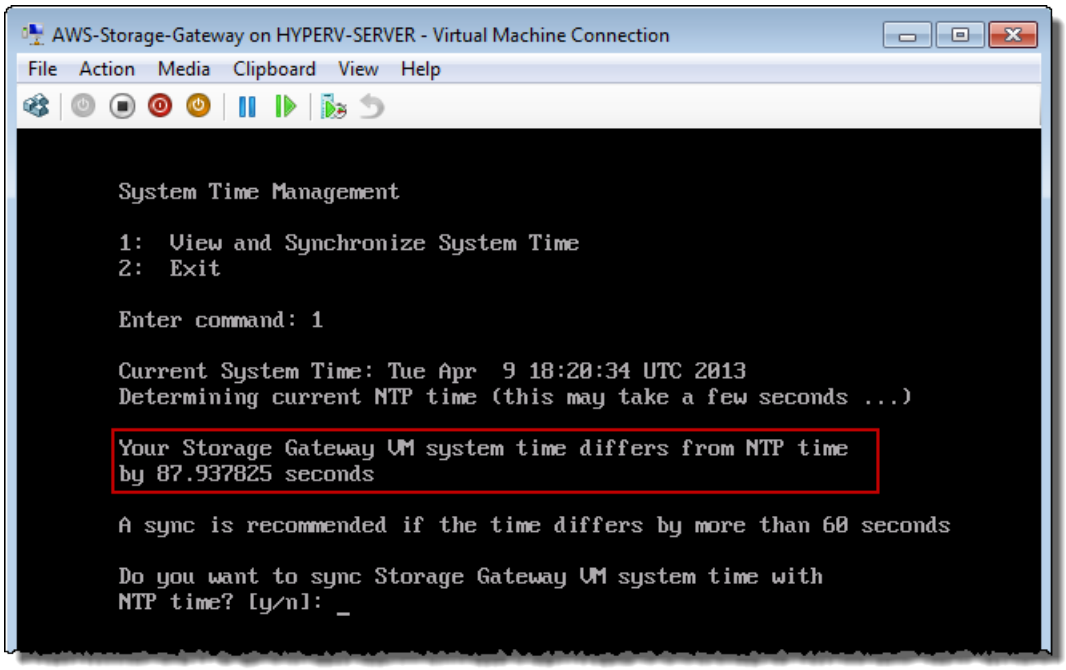
AWS-Storage-Gateway on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
System Time Management
1: View and Synchronize System Time
2: Exit
Enter command: 1
Current System Time: Tue Apr 9 18:20:34 UTC 2013
Determining current NTP time (this may take a few seconds ...)
Your Storage Gateway VM system time differs from NTP time
by 87.937825 seconds
A sync is recommended if the time differs by more than 60 seconds
Do you want to sync Storage Gateway VM system time with
NTP time? [y/n]: _

## Running Storage Gateway Commands on the Local Console

The AWS Storage Gateway console helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the console commands, you can perform maintenance tasks such as saving a routing table or connecting to AWS Support.

### To run a configuration or diagnostic command

1. Log in to your gateway's local console. For instructions, see Logging In to Your Gateway-VTL's Local Console (p. 321).
2. On the **AWS Storage Gateway Configuration** main menu, type option **5** for **Gateway Console**.

3. On the AWS Storage Gateway console, type **h**, and then press the **Return** key.

The console displays the **Available Commands** menu listing available commands and a **Gateway Console** prompt, as shown in the following screenshot.

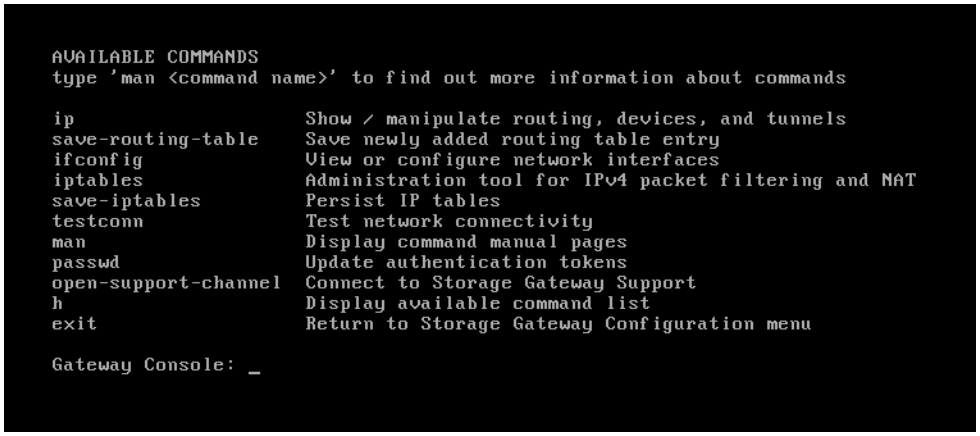
AVAILABLE COMMANDS
type 'man <command name>' to find out more information about commands
ip Show / manipulate routing, devices, and tunnels
save-routing-table Save newly added routing table entry
ifconfig View or configure network interfaces
iptables Administration tool for IPv4 packet filtering and NAT
save-iptables Persist IP tables
testconn Test network connectivity
man Display command manual pages
passwd Update authentication tokens
open-support-channel Connect to Storage Gateway Support
h Display available command list
exit Return to Storage Gateway Configuration menu
Gateway Console: _

4. To learn about a command, type **man** + ***command name*** at the **Gateway Console** prompt.

## Updating Your System

You can use the local console to update the software on your system.

### To update software on your system

1. Log in to your gateway's local console. For instructions, see Logging In to Your Gateway-VTL's Local Console (p. 321).
2. In the **AWS Storage Gateway Configuration** main menu, type option **6** to begin updating.

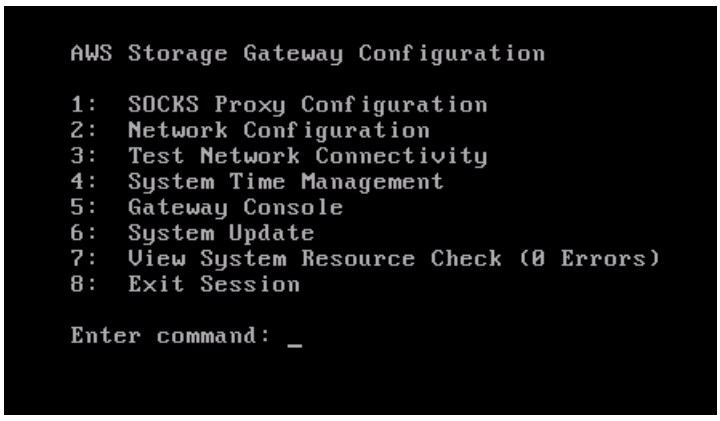

The console displays the **Apply System Updates** menu, as shown in the following screenshot.

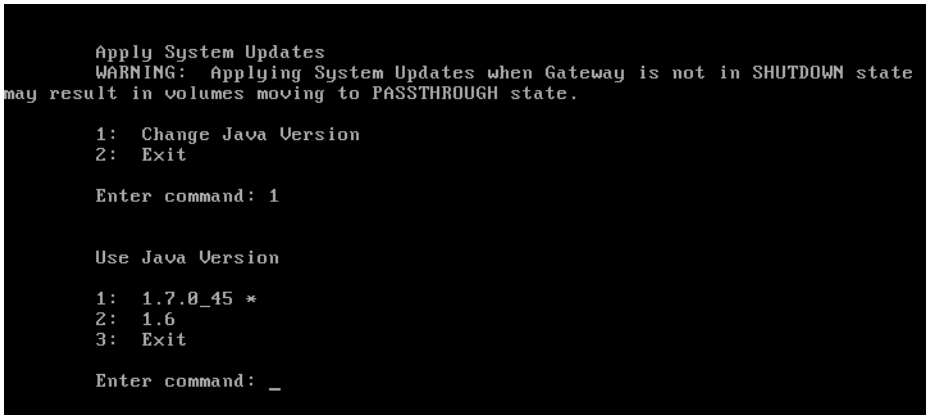
Apply System Updates
WARNING: Applying System Updates when Gateway is not in SHUTDOWN state may result in volumes moving to PASSTHROUGH state.
1: Change Java Version
2: Exit
Enter command: 1
Use Java Version
1: 1.7.0_45 *
2: 1.6
3: Exit
Enter command: _

3. Type one of the following options to use a new version of Java or to exit to the previous menu.

| To... | Do this... |
|---|---|
| Use Java version 1.7.0_45 * | Type **1**. |
| Use Java version 1.6 | Type **2**. |
| Exit to the previous menu | Type **3**. |

## Viewing Your Gateway-VTL System Resource Errors

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM and determines whether these system resources are sufficient for your gateway to function properly. You can view the results of this check on the gateway's local console.

### To view the status of a system resource check

1. Log in to your gateway-VTL's local console. For instructions, see Logging In to Your Gateway-VTL's Local Console (p. 321).
2. In the **AWS Storage Gateway Configuration** main menu, type **7** to view the results of a system resource check.

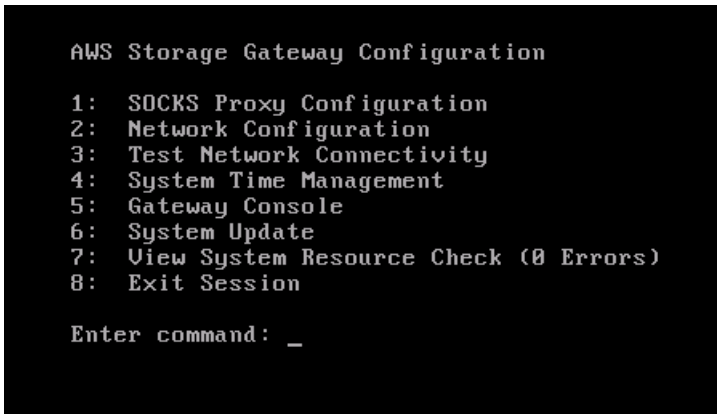

The console displays an **[OK]**, **[WARNING]**, or **[FAIL]** message for each resource as described in the table following.

| Message | Description |
|---|---|
| **[OK]** | The resource has passed the system resource check. |
| **[WARNING]** | The resource does not meet the recommended requirements, but your gateway will continue to function. AWS Storage Gateway displays a message that describes the results of the resource check. |
| **[FAIL]** | The resource does not meet the minimum requirements. Your gateway might not function properly. AWS Storage Gateway displays a message that describes the results of the resource check. |

The console also displays the number of errors and warnings next to the resource check menu option.

The following screenshot shows a **[FAIL]** message and the accompanying error message.

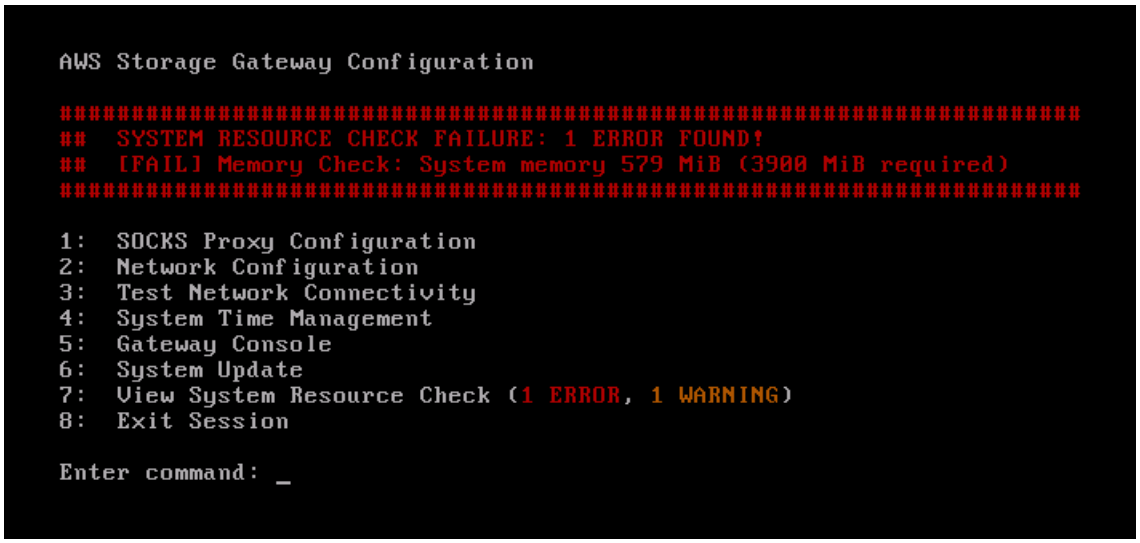
AWS Storage Gateway Configuration
##  SYSTEM RESOURCE CHECK FAILURE: 1 ERROR FOUND!
##  [FAIL] Memory Check: System memory 579 MiB (3900 MiB required)
1:  SOCKS Proxy Configuration
2:  Network Configuration
3:  Test Network Connectivity
4:  System Time Management
5:  Gateway Console
6:  System Update
7:  View System Resource Check (1 ERROR, 1 WARNING)
8:  Exit Session
Enter command: _

## Configuring Your Gateway-VTL for Multiple NICs

If you configure your gateway-VTL to use multiple network adapters (NICs), it can be accessed by more than one IP address. You might want to do this in the following situations:

- **Maximizing throughput** – You might want to maximize throughput to a gateway-VTL when network adapters are a bottleneck.
- **Application separation** – You might need to separate how your applications write to a gateway-VTL's storage volumes. For example, you might choose to have a critical storage application exclusively use one particular adapter defined for your gateway-VTL.
- **Network constraints** – Your application environment might require that you keep your iSCSI targets and the initiators that connect to them in an isolated network that is different from the network by which the gateway-VTL communicates with AWS.

In a typical multiple-adapter use case, one adapter is configured as the route by which the gateway-VTL communicates with AWS (that is, as the default gateway). Except for this one adapter, initiators must be in the same subnet as the adapter that contains the iSCSI targets to which they connect. Otherwise, communication with the intended targets might not be possible. If a target is configured on the same adapter that is used for communication with AWS, then iSCSI traffic for that target and AWS traffic will flow through the same adapter.

The following procedure assumes that your gateway-VTL VM already has one network adapter defined and that you are adding a second adapter. The first procedure shows how to add an adapter for VMware ESXi, and the second shows how for Microsoft Hyper-V.

### To configure your gateway-VTL to use an additional network adapter (VMware ESXi)

1. Shut down the gateway. For instructions, see To shut down a gateway (p. 311).
2. In the VMware vSphere client, select your gateway-VTL VM.

   The VM can remain turned on for this procedure.
3. In the client, right-click the name of your gateway-VTL VM, and click **Edit Settings**.

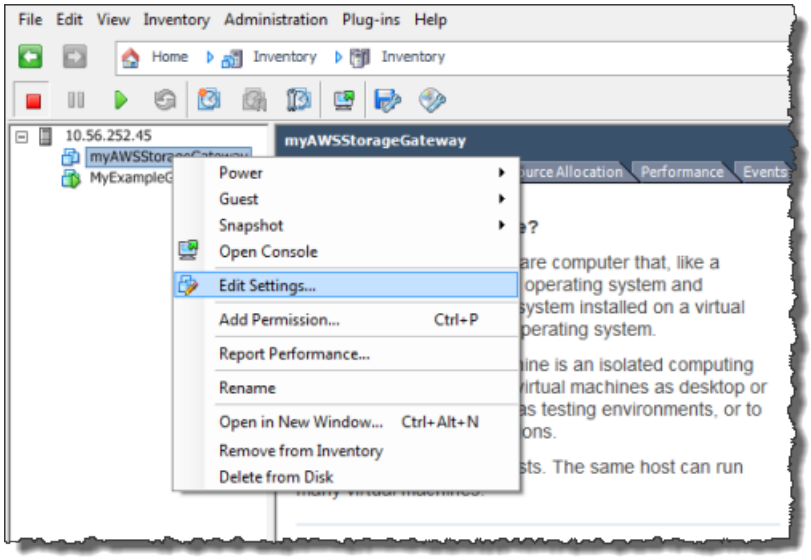


4. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add** to add a device.

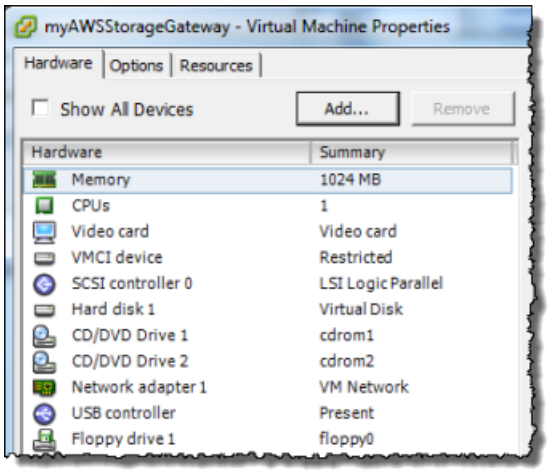

5. Follow the Add Hardware wizard to add a network adapter.

   a. In the **Device Type** pane, click **Ethernet Adapter** to add an adapter, and then click **Next**.

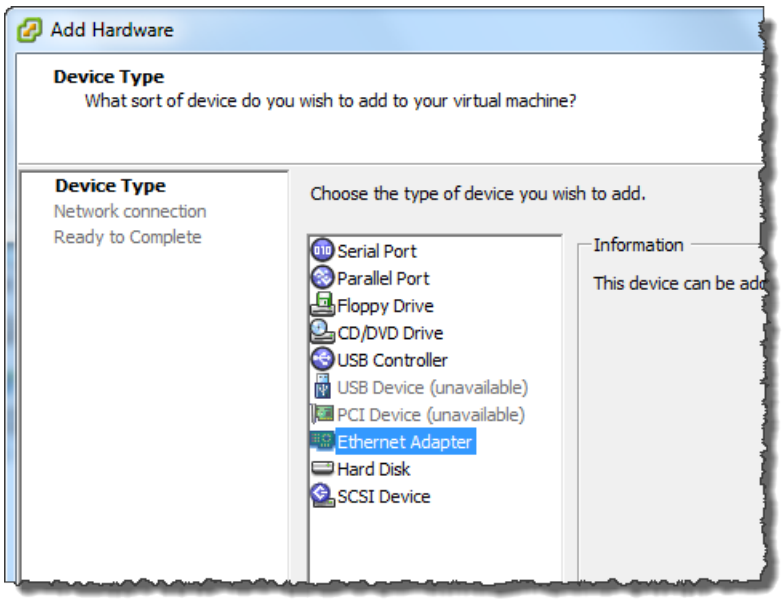


   b. In the **Type** list of the **Network Type** pane, select an adapter type, ensure that **Connect at power on** is selected, and then click **Next**.

      We recommend that you use the E1000 network adapter with AWS Storage Gateway. For more information on the adapter types that might appear in the adapter list, go to *Network Adapter Types* in the ESXi and vCenter Server Documentation.

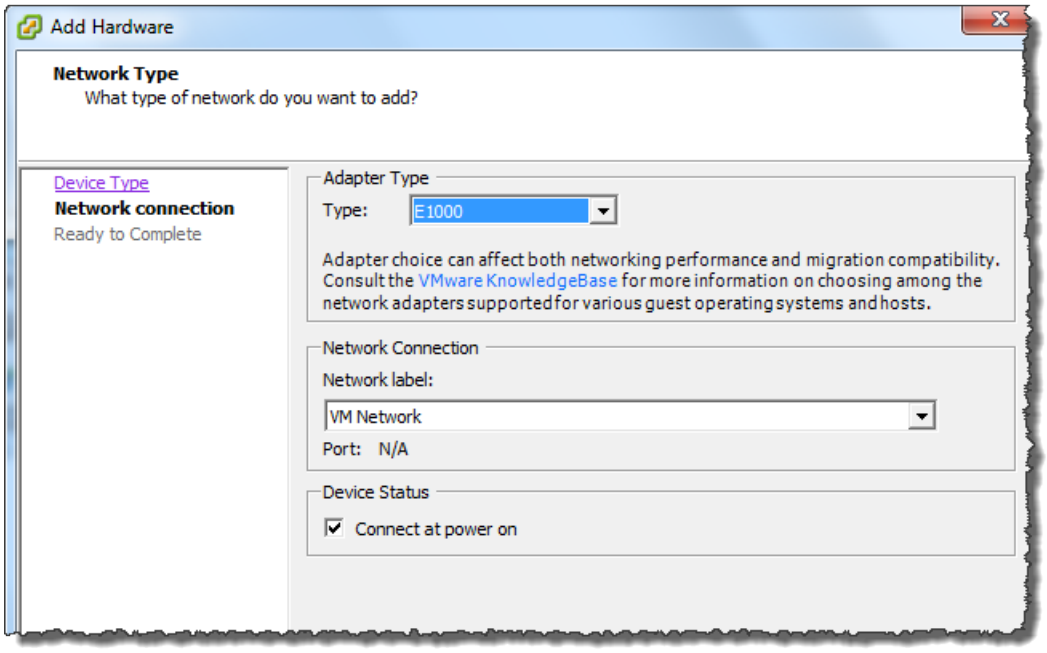


   c. In the **Ready to Complete** pane, review the information, and then click **Finish**.

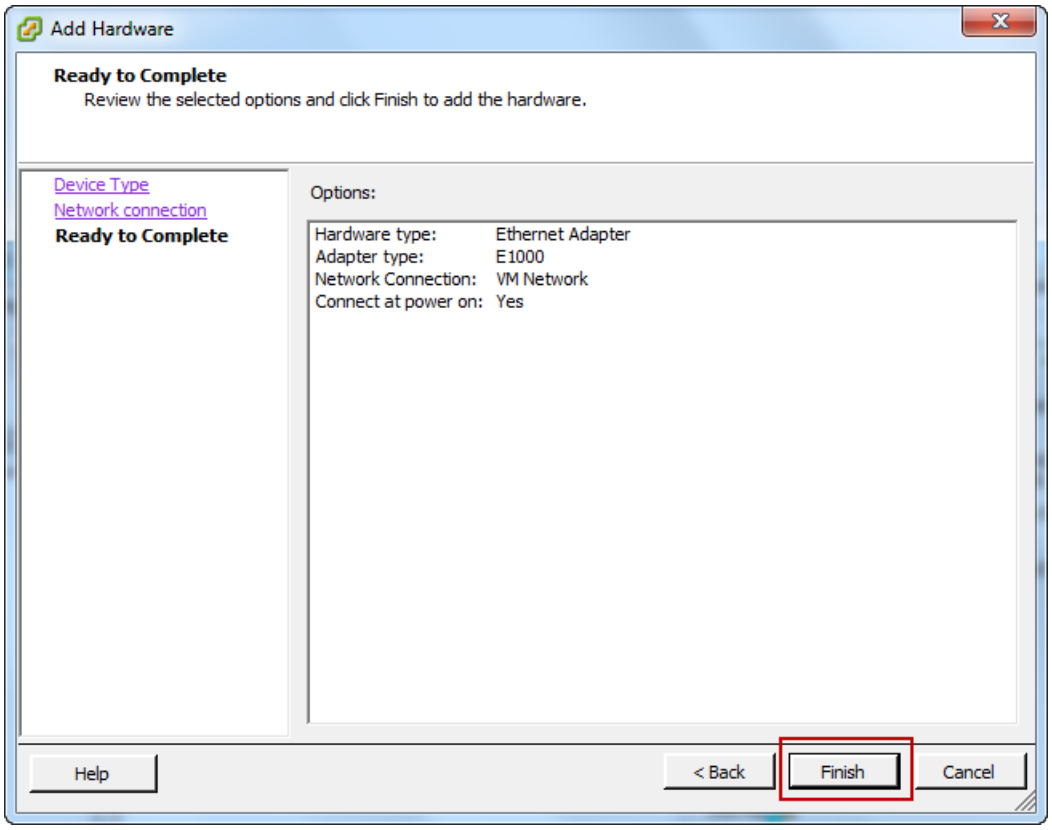


6. Click the **Summary** tab of the VM, and click **View All** next to the **IP Address** box. A **Virtual Machine IP Addresses** window displays all the IP addresses you can use to access the gateway. Confirm that a second IP address is listed for the gateway.

   **Note**
   It might take several moments for the adapter changes to take effect and the VM summary information to refresh.

   The following image is for illustration only. In practice, one of the IP addresses would be the address by which the gateway-VTL communicates to AWS and the other would be an address in a different subnet.

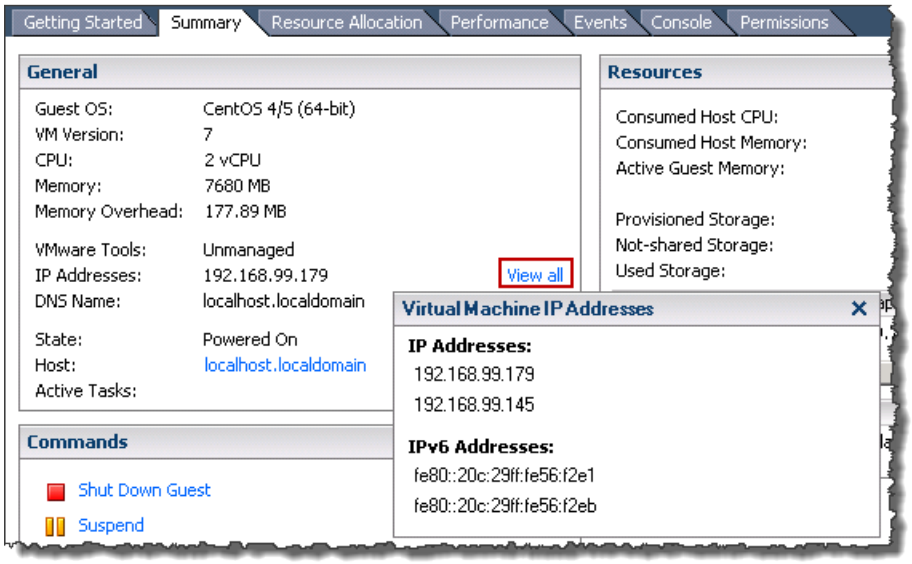


7. On the AWS Storage Gateway console, turn on the gateway. For instructions, see To start a gateway (p. 312).
8. In the **Navigation** pane of the AWS Storage Gateway console, click the gateway-VTL to which you added the adapter, and then click the **Gateway** tab.

   Confirm that the second IP address is listed in the **IP Addresses** box.

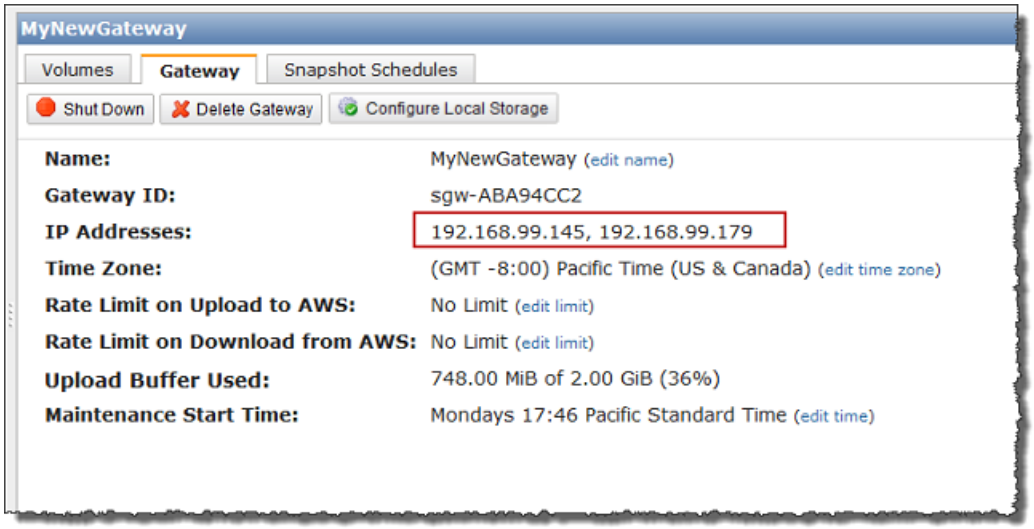
MyNewGateway
Volumes
Gateway
Snapshot Schedules
Shut Down
Delete Gateway
Configure Local Storage
Name: MyNewGateway (edit name)
Gateway ID: sgw-ABA94CC2
IP Addresses: 192.168.99.145, 192.168.99.179
Time Zone: (GMT -8:00) Pacific Time (US & Canada) (edit time zone)
Rate Limit on Upload to AWS: No Limit (edit limit)
Rate Limit on Download from AWS: No Limit (edit limit)
Upload Buffer Used: 748.00 MiB of 2.00 GiB (36%)
Maintenance Start Time: Mondays 17:46 Pacific Standard Time (edit time)

### To configure your gateway-VTL to use an additional network adapter (Microsoft Hyper-V)

1. On the AWS Storage Gateway console, turn off the gateway-VTL. For instructions, see To shut down a gateway (p. 311).
2. In the Microsoft Hyper-V Manager, select your gateway-VTL VM.
3. If the VM it isn't turned off already, right-click the gateway and select **Turn Off**.
4. In the client, right-click the name of your gateway-VTL VM and click **Settings**.

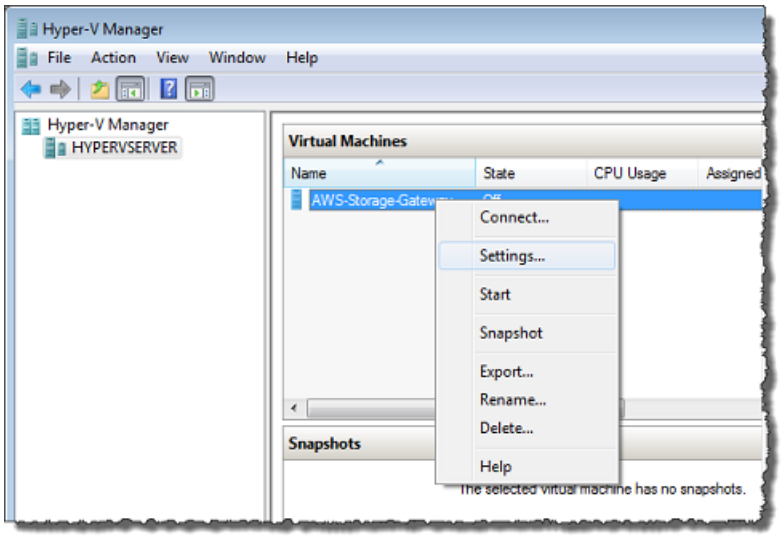


5. In the **Hardware** list in the **Settings** dialog box for the VM, click **Add Hardware**.
6. In the **Add Hardware** pane, select **Network Adapter**, and then click **Add** to add a device.

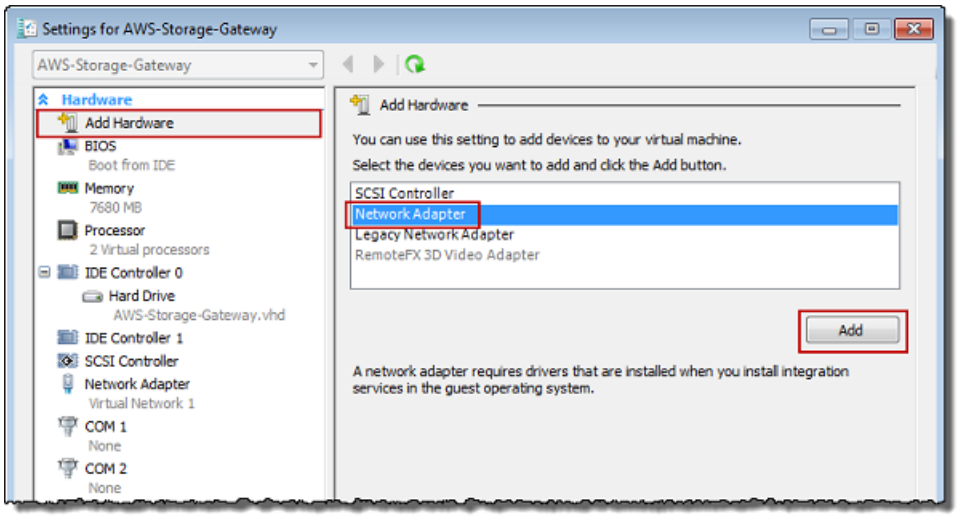


7. Configure the network adapter, and then click **Apply** to apply settings.

   In the following example, **Virtual Network 2** is selected for the new adapter.

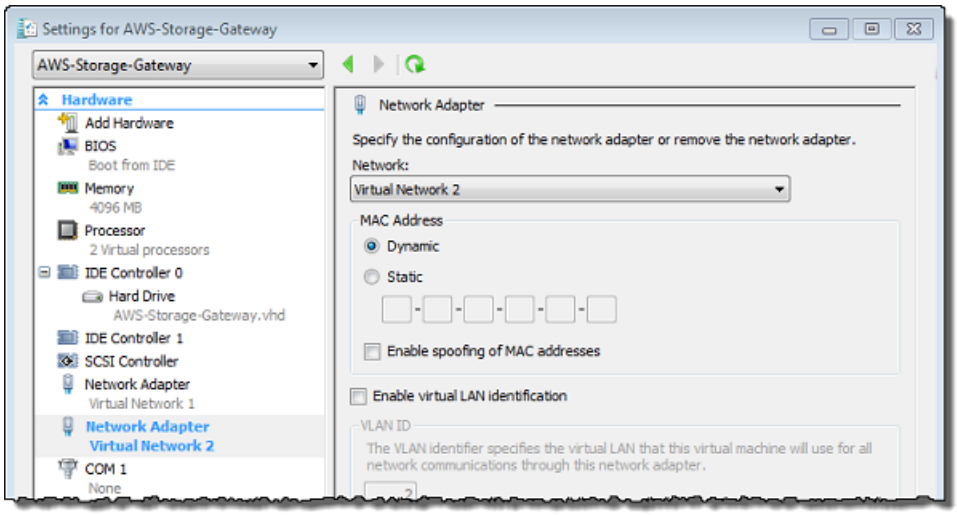


8. In the **Hardware** list of the **Settings** dialog box, confirm that the second adapter was added, and then click **OK**.
9. On the AWS Storage Gateway console, turn on the gateway. For instructions, see To start a gateway (p. 312).
10. In the **Navigation** pane of the AWS Storage Gateway console, select the gateway to which you added the adapter, and then click the **Gateway** tab.

    Confirm that the second IP address is listed in the **IP Addresses** field.

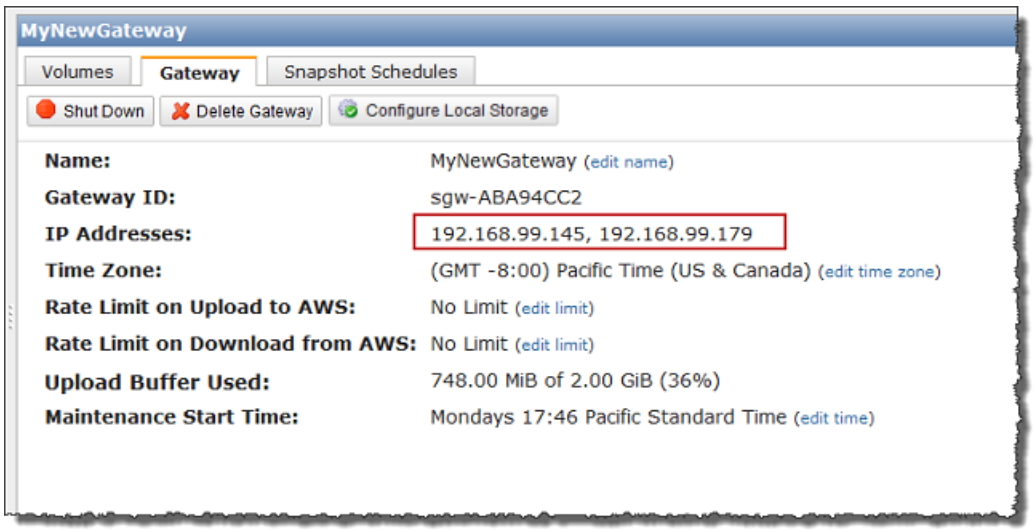
MyNewGateway
Volumes
Gateway
Snapshot Schedules
Shut Down
Delete Gateway
Configure Local Storage
Name: MyNewGateway (edit name)
Gateway ID: sgw-ABA94CC2
IP Addresses: 192.168.99.145, 192.168.99.179
Time Zone: (GMT -8:00) Pacific Time (US & Canada) (edit time zone)
Rate Limit on Upload to AWS: No Limit (edit limit)
Rate Limit on Download from AWS: No Limit (edit limit)
Upload Buffer Used: 748.00 MiB of 2.00 GiB (36%)
Maintenance Start Time: Mondays 17:46 Pacific Standard Time (edit time)

## Optimizing Gateway-VTL Performance

This section provides information about how to optimize the performance of your gateway. The guidance is based on adding resources to your gateway and adding resources to your application server.

### Add Resources to Your Gateway

**Use higher-performance disks**

To optimize gateway performance, you can add high performance disks such as serial attached SCSI (SAS) disks and solid-state drives (SSDs). You can also attach virtual disks to your VM directly from

a storage area network (SAN) instead of through the VMware Virtual Machine File System (VMFS) layer or the Microsoft Hyper-V New Technology File System (NTFS). Improved disk performance generally results in better throughput and more input/output operations per second (IOPS). To measure throughput, use the `ReadBytes` and `WriteBytes` metrics with the `Samples` Amazon CloudWatch statistic. For example, the `Samples` statistic of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you the IOPS. As a general rule, when you review these metrics for a gateway-VTL, look for low throughput and low IOPS trends to indicate disk-related bottlenecks. For more information about gateway metrics, see Measuring Performance Between Your Gateway-VTL and AWS (p. 353).

**Add CPU resources to your gateway-VTL host**

The minimum requirement for a gateway-VTL host server is four virtual processors. To optimize gateway performance, you should confirm that the four virtual processors that are assigned to the gateway VM are backed by four cores and that you are not oversubscribing the CPUs of the host server. When you add additional CPUs to your gateway-VTL host server, you increase the processing capability of the gateway-VTL to deal with, in parallel, both storing data from your application to your local storage and uploading this data to Amazon S3. Additional CPUs also help ensure that your gateway gets enough CPU resources when the host is shared with other VMs. Providing enough CPU resources has the general effect of improving throughput.

**Back gateway virtual disks with separate physical disks**

When you provision disks in a gateway-VTL setup, we strongly recommend that you do not provision local disks for the upload buffer and cache storage that use the same underlying physical storage disk. For example, for VMware ESXi, the underlying physical storage resources are represented as a data store. When you deploy the gateway-VTL VM, you choose a data store on which to store the VM files. When you provision a virtual disk (for example, to use as an upload buffer), you have the option to store the virtual disk in the same data store as the VM or a different data store. If you have more than one data store, then we strongly recommend that you choose one data store for each type of local storage you are creating. A data store that is backed by only one underlying physical disk, or that is backed by a less-performant RAID configuration such as RAID 1, can lead to poor performance—for example, when used to back both the cache storage and upload buffer in a gateway-VTL setup.

### Add Resources to Your Application Environment

**Increase the bandwidth between your application server and your gateway**-VTL

To optimize gateway performance, ensure that the network bandwidth between your application and the gateway-VTL can sustain your application needs. You can use the `ReadBytes` and `WriteBytes` metrics of the gateway to measure the total data throughput (for more information on these metrics, see Measuring Performance Between Your Gateway-VTL and AWS (p. 353)). For your application, compare the measured throughput with the desired throughput. If the measured throughput is less than the desired throughput, then increasing the bandwidth between your application and gateway-VTL can improve performance if the network is the bottleneck. Similarly, you can increase the bandwidth between your VM and your local disks, if they're not direct-attached.

**Add CPU resources to your application environment**

If your application can make use of additional CPU resources, then adding more CPUs can help your application to scale its I/O load.

## Maintenance Tasks When Running a Gateway-VTL on an Amazon EC2 Instance

This section describes maintenance tasks for a gateway-VTL deployed on Amazon Elastic Compute Cloud (Amazon EC2). For information about managing a gateway-VTL deployed on-premises, see Maintenance Tasks When Running a Gateway-VTL On-Premises (p. 321).

| To... | Do This |
|---|---|
| Manage cache storage and upload buffer space by adding or removing Amazon EBS volumes for your Amazon EC2 instance. | See Adding and Removing Amazon EBS Volumes for Your Gateway-VTL Instance (p. 343). |
| Enable gateway-VTL access to AWS Support to investigate and troubleshoot issues, and then disable access as needed. | See Enabling AWS Support to Access a Gateway Hosted on an Amazon EC2 Instance (p. 369). |
| Configure a security group to set gateway-required ports or to allow access to your Amazon EC2 instance. | See Configuring Security Groups for Your Amazon EC2 Gateway-VTL Instance (p. 344). |
| Free AWS resources so you don't incur charges, if you don't plan to continue using your gateway-VTL. | See Delete an Amazon EC2 Gateway-VTL and Associated Resources (p. 344). |

## Adding and Removing Amazon EBS Volumes for Your Gateway-VTL Instance

When you initially configured your gateway-VTL to run as an Amazon EC2 instance, you allocated Amazon EBS volumes for use as an upload buffer and cache storage. Over time, as your applications needs change, you can allocate additional Amazon EBS volumes for this use. You can also reduce the storage you allocated by removing previously allocated Amazon EBS volumes. For more information about Amazon EBS, go to Amazon Elastic Block Store (Amazon EBS) in the *Amazon EC2 User Guide for Linux Instances*.

Before you add more storage to the gateway, you should review how to size your upload buffer and cache storage based on your application needs for a gateway-VTL. To do so, see Sizing the Upload Buffer (p. 305).

There is a limit to the maximum storage you can allocate as an upload buffer and cache storage. You can attach as many Amazon EBS volumes to your instance as you want, but you can only configure these volumes as upload buffer and cache storage space up to these storage limits. For more information, see Appendix F: AWS Storage Gateway Limits (p. 421).

### To create an Amazon EBS volume, attach it, and configure it for your gateway-VTL

1. Create an Amazon EBS volume. For instructions, go to Creating or Restoring an Amazon EBS Volume in the *Amazon EC2 User Guide for Linux Instances*.
2. Attach the Amazon EBS volume to your Amazon EC2 instance. For instructions, go to Attaching an Amazon EBS Volume to an Instance in the *Amazon EC2 User Guide for Linux Instances*.
3. Configure the Amazon EBS volume you added as either an upload buffer or cache storage. For instructions, see Configuring an Upload Buffer and Cache Storage (Gateway-VTL) (p. 310).

There are times you might find you don't need the amount of storage you allocated for the upload buffer.

### To remove an Amazon EBS volume

**Warning**
These steps apply only for Amazon EBS volumes allocated as upload buffer space. If you remove an Amazon EBS volume that is allocated as cache storage from a gateway-VTL, virtual tapes on the gateway-VTL will have the IRRECOVERABLE status, and you risk data loss. For more information on the IRRECOVERABLE status, see Tape Status Information (p. 278).

1. Shut down the gateway-VTL by following the approach described in the Shutting Down and Starting Your Gateway-VTL (p. 311) section.
2. Detach the Amazon EBS volume from your Amazon EC2 instance. For instructions, go to Detaching an Amazon EBS Volume from an Instance in the *Amazon EC2 User Guide for Linux Instances*.

3. Delete the Amazon EBS volume. For instructions, go to Deleting an Amazon EBS Volume in the *Amazon EC2 User Guide for Linux Instances*.
4. Start the gateway-VTL by following the approach described in the Shutting Down and Starting Your Gateway-VTL (p. 311) section.

## Configuring Security Groups for Your Amazon EC2 Gateway-VTL Instance

A security group controls traffic to your Amazon EC2 gateway-VTL instance. When you create an instance from the Amazon Machine Image (AMI) for AWS Storage Gateway from AWS Marketplace, you have two choices for launching the instance. The Step 2: Deploy and Activate a Gateway on Amazon EC2 (p. 58) section shows how to launch the instance by using the **Launch with an EC2 Console** feature of AWS Marketplace. Launching this way is the approach we recommend. You can also launch an instance by using the **1-Click Launch** feature in AWS Marketplace. In this case, an autogenerated security group that is named `AWS Storage Gateway-1-0-AutogenByAWSMP-` is created. This security group has the correct rule for port 80 to let you activate your gateway. For more information about security groups, go to Security Group Concepts in the *Amazon EC2 User Guide for Linux Instances*.

Regardless of the security group that you use, we recommend that:

- The security group should not allow incoming connections from the outside Internet. It should allow only instances within the gateway-VTL's security group to communicate with the gateway-VTL. If you need to allow instances to connect to the gateway-VTL from outside its security group, we recommend that you allow connections only on ports 3260 (for iSCSI connections) and 80 (for activation).
- If you want to activate your gateway-VTL from a host outside the gateway-VTL's security group, allow incoming connections on port 80 from the IP address of that host. If you cannot determine the activating host's IP address, you can open port 80, activate your gateway, and then close access on port 80 after completing activation.
- You allow port 22 access only if you are using AWS Support for troubleshooting purposes. For more information, see Enabling AWS Support to Access a Gateway Hosted on an Amazon EC2 Instance (p. 369).

If you are using an Amazon EC2 instance as an initiator (that is, to connect to the iSCSI targets on the gateway-VTL you deployed on Amazon EC2), then we recommend a two-step approach.

1. You should put the initiator instance in the same security group as the gateway-VTL.
2. You should configure access so the initiator can communicate with the gateway-VTL.

## Delete an Amazon EC2 Gateway-VTL and Associated Resources

If you don't plan to continue using your Amazon EC2 gateway-VTL, consider deleting the gateway-VTL and its associated resources. Removing unnecessary resources helps reduce your monthly bill. To delete the gateway-VTL and its associated resources, we suggest the following sequence of actions.

1. Delete all virtual tapes that have RETRIEVING status from the gateway-VTL. For more information, see Canceling Tape Retrieval (p. 281).
2. Delete all virtual tapes you have retrieved to your gateway-VTL. For more information, see Deleting Tapes from Your Gateway-VTL (p. 280).
3. Delete all virtual tapes from the tape library. For more information, see Deleting Tapes from Your Gateway-VTL (p. 280).
4. Delete the gateway-VTL itself. For more information, see Deleting Your Gateway-VTL (p. 320).
5. Terminate all Amazon EC2 instances, and delete all Amazon EBS volumes. For more information, see Clean Up Your Instance and Volume in the *Amazon EC2 User Guide for Linux Instances*.
6. Delete all archived virtual tapes from the virtual tape shelf (VTS). For more information, see Deleting Tapes from the VTS (p. 288).

**Important**
If you retrieve a virtual tape that has not been stored in the VTS for a least 90 days, you are charged a minimum of 90-days storage.

## Performing Maintenance Tasks on the Local Console for an Amazon EC2 Gateway-VTL

Some maintenance tasks require that you log in to the local console when running a gateway-VTL deployed on an Amazon EC2 instance. This section describes how to log into the local console and perform maintenance tasks.

**Topics**


- Logging In to the Local Console for an Amazon EC2 Gateway-VTL (p. 345)
- Routing Your Gateway-VTL Through a Proxy (p. 346)
- Testing Your Gateway-VTL's Connectivity to the Internet (p. 347)
- Running Storage Gateway Commands on the Local Console (p. 348)
- Updating Your System (p. 349)
- Viewing Your Gateway-VTL System Resource Errors (p. 350)


### Logging In to the Local Console for an Amazon EC2 Gateway-VTL

You can connect to your Amazon EC2 instance by using a Secure Shell (SSH) client. For detailed information, go to Connect to Your Instance in the *Amazon EC2 User Guide*. To connect this way, you will need the SSH key pair you specified when you launched the instance. For information about Amazon EC2 key pairs, go to Amazon EC2 Key Pairs in the *Amazon EC2 User Guide.*

#### To log in to the gateway-VTL's local console

- Log in to your local console. If you are connecting from a Windows computer, log in as *sguser*.

After you log in, you will see the **AWS Storage Gateway Configuration** main menu, as shown in the following screenshot.

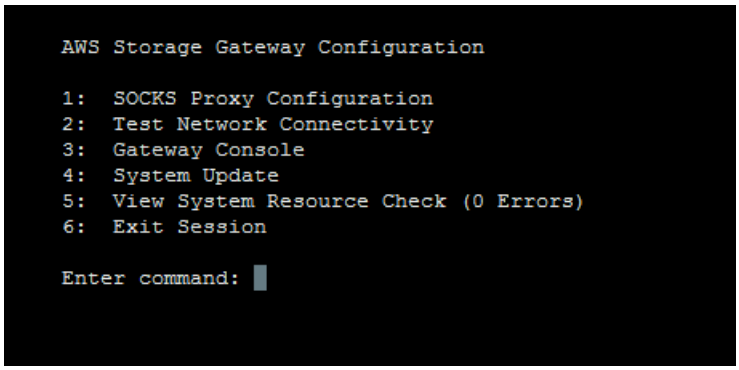


You use the **AWS Storage Gateway Configuration** menus to perform configuration tasks.

| To... | See... |
|---|---|
| Configure a SOCKS proxy for your gateway-VTL | Routing Your Gateway-VTL Through a Proxy (p. 325). |

| To... | See... |
|---|---|
| Test network connectivity | Testing Your Gateway-VTL's Connectivity to the Internet (p. 329) |
| Run Storage Gateway console commands | Running Storage Gateway Commands on the Local Console (p. 332) |
| Update your system | Updating Your System (p. 333) |

To exit the configuration session, type **5** to exit the menu.

## Routing Your Gateway-VTL Through a Proxy

AWS Storage Gateway supports the configuration of a Socket Secure version 5 (SOCKS5) proxy between your gateway-VTL and AWS.

> **Note**
> The only proxy configuration AWS Storage Gateway supports is SOCKS5.

If your gateway-VTL must use a proxy server to communicate to the Internet, then you need to configure the SOCKS proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy. After you do this, AWS Storage Gateway will route all HyperText Transfer Protocol Secure (HTTPS) traffic through your proxy server.

### To route your gateway-VTL Internet traffic through a local proxy server

1. Log in to your gateway's local console. For instructions, see Logging In to the Local Console for an Amazon EC2 Gateway-VTL (p. 345).
2. In the **AWS Storage Gateway Configuration** main menu, type **1** to begin configuring the SOCKS proxy.

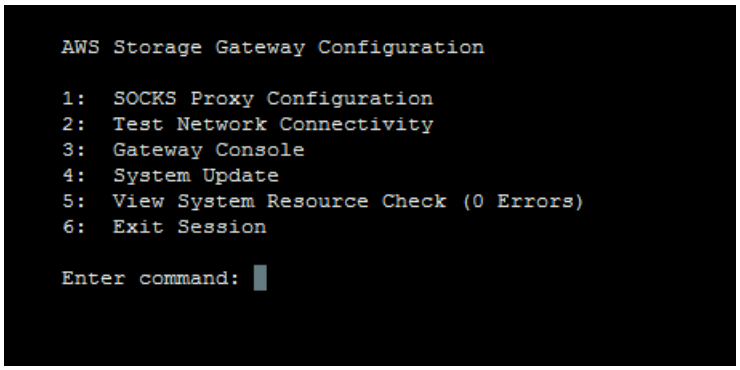


3. Choose one of the following options in the **AWS Storage Gateway SOCKS Proxy Configuration** menu:

AWS Storage Gateway SOCKS Proxy Configuration
1: Configure SOCKS Proxy
2: View Current SOCKS Proxy Configuration
3: Remove SOCKS Proxy Configuration
4: Exit
Enter command:

| To... | Do this... |
|---|---|
| Configure a SOCKS proxy | Type option **1**.<br><br>You will need to supply a host name and port to complete configuration. |
| View the current SOCKS proxy configuration | Type option **2**.<br><br>If a SOCKS proxy is not configured, the message `SOCKS Proxy not configured` is displayed. If a SOCKS proxy is configured, the host name and port of the proxy are displayed. |
| Remove a SOCKS proxy configuration | Type option **3**.<br><br>The message `SOCKS Proxy Configuration Removed` is displayed. |
| Exit this menu and return to the previous menu | Type option **4**. |

## Testing Your Gateway-VTL's Connectivity to the Internet

You can use your gateway's local console to test your Internet connection. This test can be useful when you are troubleshooting network issues with your gateway-VTL. The network connectivity test does not take into account any SOCKS proxy you might have configured.

### To test your gateway-VTL's connection to the Internet

1. Log in to your gateway's local console. For instructions, see Logging In to the Local Console for an Amazon EC2 Gateway-VTL (p. 345).
2. In the **AWS Storage Gateway Configuration** main menu, type option **2** to begin testing network connectivity.

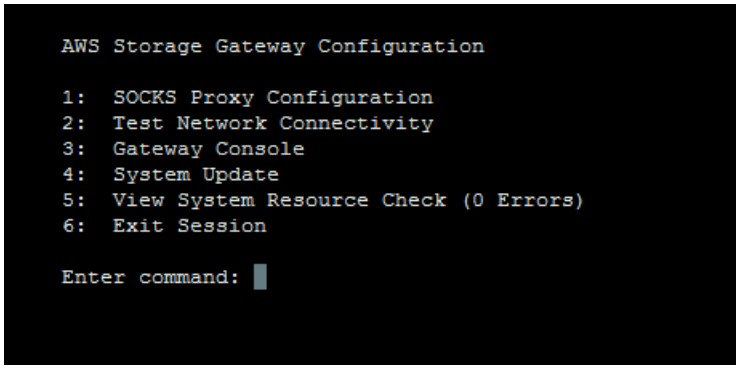


   The console displays available regions.
3. Select the region you want to test, as shown in the following screenshot.

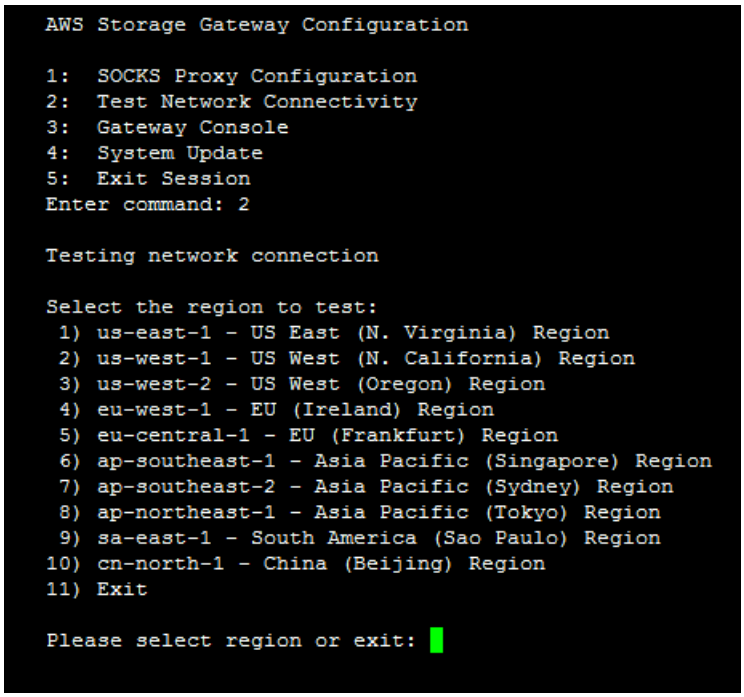
AWS Storage Gateway Configuration
1: SOCKS Proxy Configuration
2: Test Network Connectivity
3: Gateway Console
4: System Update
5: Exit Session
Enter command: 2
Testing network connection
Select the region to test:
1) us-east-1 - US East (N. Virginia) Region
2) us-west-1 - US West (N. California) Region
3) us-west-2 - US West (Oregon) Region
4) eu-west-1 - EU (Ireland) Region
5) eu-central-1 - EU (Frankfurt) Region
6) ap-southeast-1 - Asia Pacific (Singapore) Region
7) ap-southeast-2 - Asia Pacific (Sydney) Region
8) ap-northeast-1 - Asia Pacific (Tokyo) Region
9) sa-east-1 - South America (Sao Paulo) Region
10) cn-north-1 - China (Beijing) Region
11) Exit
Please select region or exit:

Each endpoint in the region you select displays either a PASSED or FAILED message, as shown following.

| Message | Description |
|---|---|
| [ PASSED ] | AWS Storage Gateway has Internet connectivity. |
| [ FAILED ] | AWS Storage Gateway does not have Internet connectivity. |

## Running Storage Gateway Commands on the Local Console

The AWS Storage Gateway console helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the console commands, you can perform maintenance tasks such as saving a routing table or connecting to AWS Support.

### To run a configuration or diagnostic command

1. Log in to your gateway's local console. For instructions, see Logging In to the Local Console for an Amazon EC2 Gateway-VTL (p. 345).
2. In the **AWS Storage Gateway Configuration** main menu, type option **3** for **Gateway Console**.

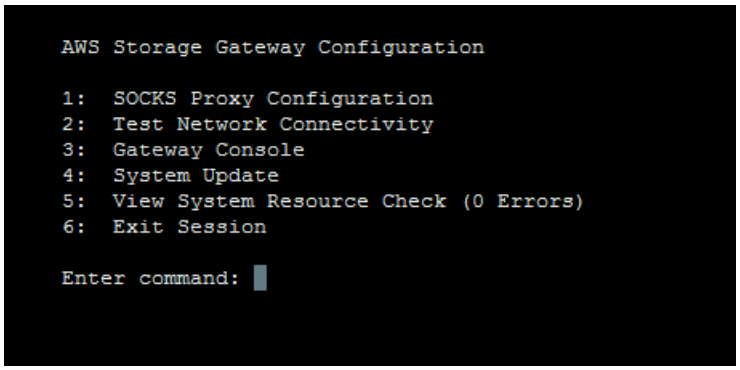


3. In the Storage Gateway console, enter **h**, and then press the **Return** key.

The console displays the **Available Commands** menu listing available commands and a **Gateway Console** prompt, as shown in the following screenshot.

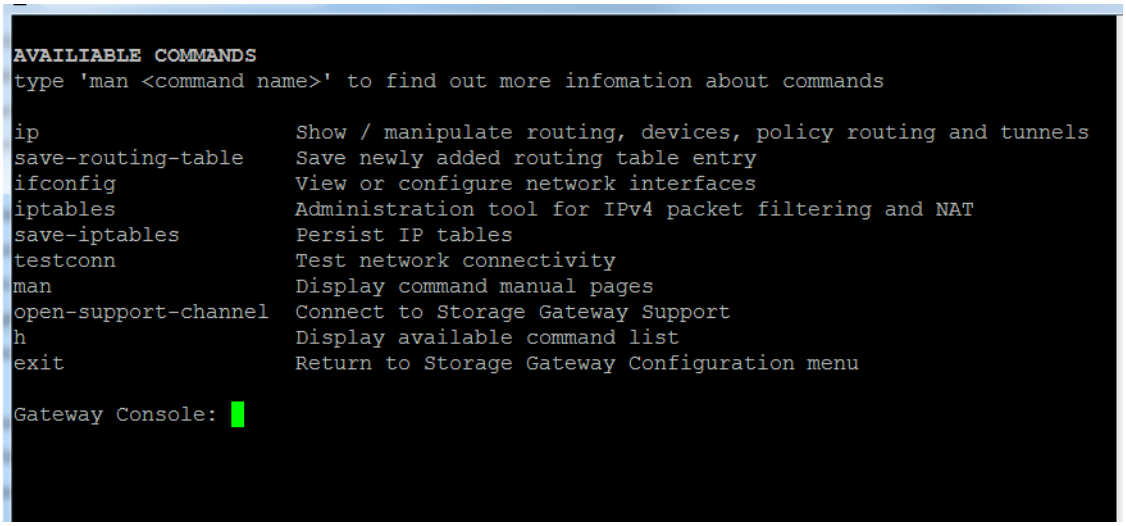
AVAILIABLE COMMANDS
type 'man <command name>' to find out more infomation about commands
ip Show / manipulate routing, devices, policy routing and tunnels
save-routing-table Save newly added routing table entry
ifconfig View or configure network interfaces
iptables Administration tool for IPv4 packet filtering and NAT
save-iptables Persist IP tables
testconn Test network connectivity
man Display command manual pages
open-support-channel Connect to Storage Gateway Support
h Display available command list
exit Return to Storage Gateway Configuration menu
Gateway Console:

4. To learn about a command, type **man** + ***command name*** at the **Gateway Console** prompt.

## Updating Your System

You can use the local console to update the software on your system.

### To update software on your system

1. Log in to your gateway's local console. For instructions, see Logging In to the Local Console for an Amazon EC2 Gateway-VTL (p. 345).
2. In the **AWS Storage Gateway Configuration** main menu, type option **4** to begin updating.

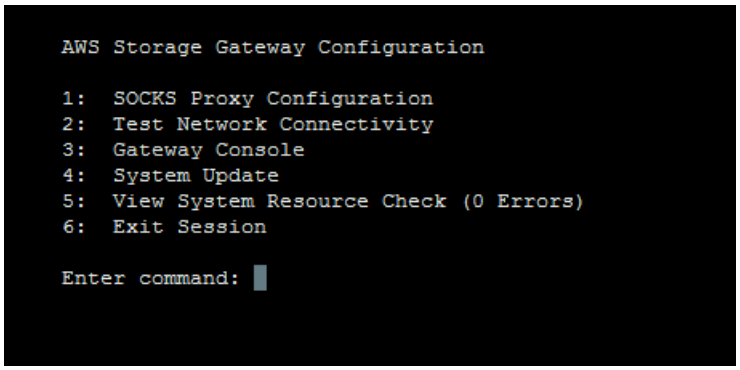


The console displays the **Apply System Updates** menu, as shown in the following screenshot.

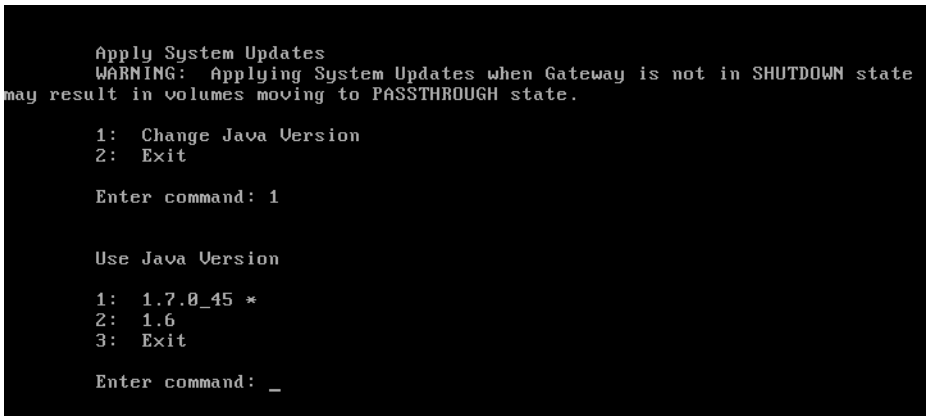
Apply System Updates
WARNING: Applying System Updates when Gateway is not in SHUTDOWN state may result in volumes moving to PASSTHROUGH state.
1: Change Java Version
2: Exit
Enter command: 1
Use Java Version
1: 1.7.0_45 *
2: 1.6
3: Exit
Enter command: _

3. Type one of the following options to use a new version of Java or to exit to the previous menu.

| To... | Do this... |
|---|---|
| Use Java version 1.6 | Type **1**. |
| Use Java version 1.7.0_45 * | Type **2**. |
| Exit to the previous menu | Type **3**. |

## Viewing Your Gateway-VTL System Resource Errors

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM and determines whether these system resources are sufficient for your gateway to function properly. You can view the results of this check on the gateway's local console.

### To view the status of a system resource check

1. Log in to your gateway-VTL's local console. For instructions, see Logging In to Your Gateway-VTL's Local Console (p. 321).

2. In the **AWS Storage Gateway Configuration** main menu, type **5** to view the results of a system resource check.

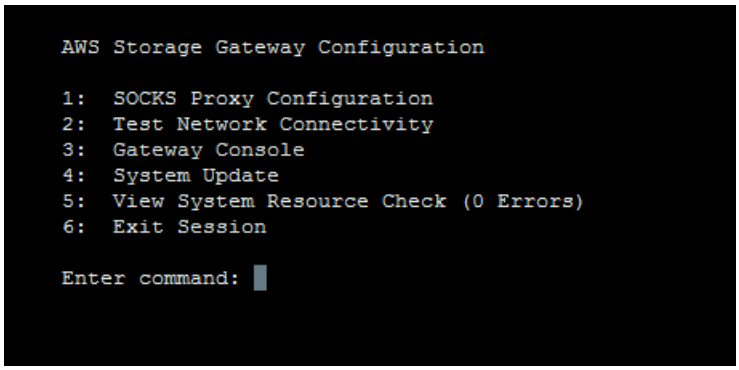


The console displays an **[OK]**, **[WARNING]**, or **[FAIL]** message for each resource as described in the table following.

| Message | Description |
|---|---|
| **[OK]** | The resource has passed the system resource check. |
| **[WARNING]** | The resource does not meet the recommended requirements, but your gateway will continue to function. AWS Storage Gateway displays a message that describes the results of the resource check. |
| **[FAIL]** | The resource does not meet the minimum requirements. Your gateway might not function properly. AWS Storage Gateway displays a message that describes the results of the resource check. |

The console also displays the number of errors and warnings next to the resource check menu option.

The following screenshot shows a **[FAIL]** message and the accompanying error message.

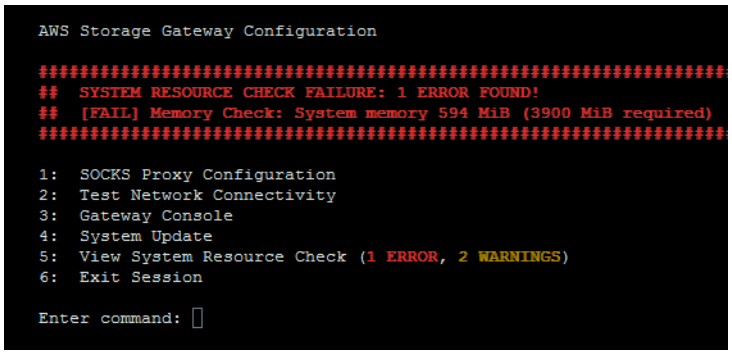
AWS Storage Gateway Configuration
## SYSTEM RESOURCE CHECK FAILURE: 1 ERROR FOUND!
## [FAIL] Memory Check: System memory 594 MiB (3900 MiB required)
1: SOCKS Proxy Configuration
2: Test Network Connectivity
3: Gateway Console
4: System Update
5: View System Resource Check (1 ERROR, 2 WARNINGS)
6: Exit Session
Enter command:

# Monitoring Your Gateway-VTL

In this section, we discuss how to monitor your gateway-VTL, virtual tapes associated with your gateway-VTL, cache storage, and the upload buffer. You use the AWS Management Console to view

metrics for your gateway-VTL. With metrics, you can track the health of your gateway-VTL and set up alarms to notify you when one or more metrics are outside a defined threshold.

AWS Storage Gateway provides Amazon CloudWatch metrics at no additional charge. AWS Storage Gateway metrics are recorded for a period of two weeks. This functionality lets you access historical information and provides you with a better perspective of how your gateway-VTL and virtual tapes are performing. For detailed information about Amazon CloudWatch, go to the Amazon CloudWatch Developer Guide.

**Topics**


- Using Amazon CloudWatch Metrics (p. 352)
- Measuring Performance Between Your Gateway-VTL and AWS (p. 353)
- Monitoring the Upload Buffer (p. 355)
- Monitoring Cache Storage (p. 358)
- Understanding Metrics (p. 359)


# Using Amazon CloudWatch Metrics

You can get monitoring data for your gateway-VTL by using either the AWS Management Console or the Amazon CloudWatch API. The console displays a series of graphs based on the raw data from the Amazon CloudWatch API. The Amazon CloudWatch API can also be used through one of the Amazon AWS Software Development Kits (SDKs) or the Amazon CloudWatch API tools. Depending on your needs, you might prefer to use either the graphs displayed in the console or retrieved from the API.

Regardless of which method you choose to use to work with metrics, you must specify the following information.

- First, you specify the metric dimension to work with. A *dimension* is a name-value pair that helps you to uniquely identify a metric. The dimensions for AWS Storage Gateway are `GatewayId` and `GatewayName`. In the Amazon CloudWatch console, the `Gateway Metrics` view is provided to easily select gateway-specific and tape-specific dimensions. For more information about dimensions, see Dimensions in the *Amazon CloudWatch Developer Guide*.
- Second, you specify the metric name, such as `ReadBytes`.

The following table summarizes the types of AWS Storage Gateway metric data that are available to you.

| Amazon Cloud-Watch Namespace | Dimension | Description |
|---|---|---|
| `AWS/StorageG-ateway` | `GatewayId`, `Gate-wayName` | These dimensions filter for metric data that describes aspects of the gateway-VTL. You can identify a gateway-VTL to work with by specifying both the `GatewayId` and the `GatewayName` dimensions.<br><br>Throughput and latency data of a gateway-VTL is based on all the virtual tapes in the gateway-VTL.<br><br>Data is available automatically in 5-minute periods at no charge. |

Working with gateway and tape metrics is similar to working with other service metrics. The Amazon CloudWatch documentation outlines many of the common tasks, listed following for your convenience:

- Listing Available Metrics
- Getting Statistics for a Metric
- Creating CloudWatch Alarms

# Measuring Performance Between Your Gateway-VTL and AWS

Data throughput, data latency, and operations per second are measures that you can use to understand how your application storage that is using your gateway-VTL is performing. When you use the correct aggregation statistic, these values can be measured by using the AWS Storage Gateway metrics that are provided for you.

A *statistic* is an aggregation of a metric over a specified period of time. When you view the values of a metric in Amazon CloudWatch, use the `Average` statistic for data latency (milliseconds), and use the `Samples` statistic for input/output operations per second (IOPS). For more information, see Statistics in the *Amazon CloudWatch Developer Guide*.

The following table summarizes the metrics and the corresponding statistic you can use to measure the throughput, latency, and IOPS between your gateway-VTL and AWS.

| Item of Interest | How to Measure |
|---|---|
| Latency | Use the `ReadTime` and `WriteTime` metrics with the `Average` Amazon CloudWatch statistic. For example, the `Average` value of the `ReadTime` metric gives you the latency per operation over the sample period of time. |
| Throughput to AWS | Use the `CloudBytesDownloaded` and `CloudBytesUploaded` metrics with the `Sum` Amazon CloudWatch statistic. For example, the `Sum` value of the `CloudBytesDownloaded` metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput from AWS to the gateway-VTL as a rate in bytes per second. |
| Latency of data to AWS | Use the `CloudDownloadLatency` metric with the `Average` statistic. For example, the `Average` statistic of the `CloudDownloadLatency` metric gives you the latency per operation. |

The following procedures assume that you start in the **Metrics** tab on the Amazon CloudWatch console.

### To measure the upload data throughput from a gateway-VTL to AWS

1. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway-VTL that you want to work with.
2. Select the `CloudBytesUploaded` metric.
3. Select a **Time Range** value.
4. Select the `Sum` statistic.
5. Select a **Period** value of 5 minutes or greater.
6. In the resulting time-ordered set of data points, divide each data point by the period (in seconds) to get the throughput at that sample period.

The following image shows the `CloudBytesUploaded` metric for a gateway tape with the `Sum` statistic. In the image, placing the cursor over a data point displays information about the data point, including its value and the number of bytes uploaded. Divide this value by the **Period** (5 minutes) to get the throughput

at that sample point. For the point highlighted, the throughput from the gateway-VTL to AWS is 555,544,576 bytes divided by 300 seconds, which is 1.7 megabytes per second.

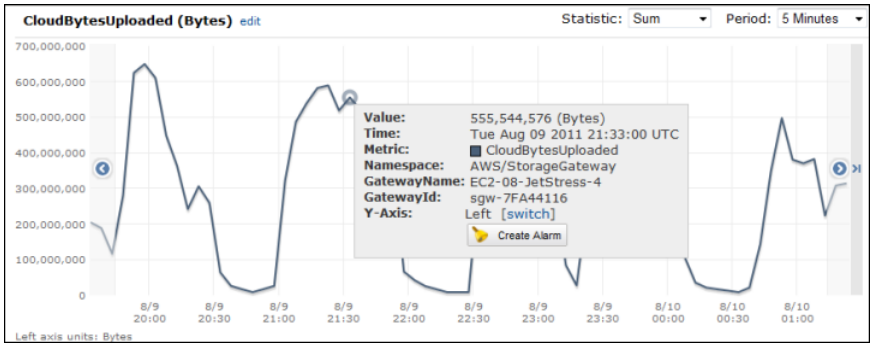


### To measure the data latency from a gateway-VTL to AWS

1. Select the **StorageGateway: GatewayMetrics** dimension, and find the gateway-VTL that you want to work with.
2. Select the `CloudDownloadLatency` metric.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the latency in milliseconds.

### To set an upper threshold alarm for a gateway-VTL's throughput to AWS

1. Start the Create Alarm Wizard.
2. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway-VTL that you want to work with.
3. Select the `CloudBytesUploaded` metric.
4. Define the alarm by defining the alarm state when the `CloudBytesUploaded` metric is greater than or equal to a specified value for a specified time. For example, you can define an alarm state when the `CloudBytesUploaded` metric is greater than 10 megabytes for 60 minutes.
5. Configure the actions to take for the alarm state. For example, you can have an email notification sent to you.
6. Create the alarm. You can create the alarm from the Amazon CloudWatch console.

### To set an upper threshold alarm for reading data from AWS

1. Start the Create Alarm Wizard.
2. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway-VTL that you want to work with.
3. Select the `CloudDownloadLatency` metric.
4. Define the alarm by defining the alarm state when the `CloudDownloadLatency` metric is greater than or equal to a specified value for a specified time. For example, you can define an alarm state when the `CloudDownloadLatency` is greater than 60,000 milliseconds for greater than 2 hours.
5. Configure the actions to take for the alarm state. For example, you can have an email notification sent to you.
6. Create the alarm. You can create the alarm from the Amazon CloudWatch console.

# Monitoring the Upload Buffer

This section discusses how to monitor a gateway-VTL's upload buffer and how to create an alarm to get notification when the buffer exceeds a specified threshold. This approach lets you proactively add buffer storage to a gateway-VTL before the buffer fills completely and your storage application stops backing up to AWS.

| Item of Interest | How to Measure |
|---|---|
| Upload buffer usage | Use the `UploadBufferPercentUsed`, `UploadBufferUsed`, and `UploadBufferFree` metrics with the `Average` statistic. For example, use the `UploadBufferUsed` metric with the `Average` statistic to analyze the storage usage over a time period. |

The following procedures assume that you start in the Amazon CloudWatch console.

### To measure upload buffer percent used

1. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway-VTL that you want to work with.
2. Select the `UploadBufferPercentUsed` metric.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percent used of the upload buffer.

The following procedure shows you how to create an alarm using the Amazon CloudWatch console and the Create Alarm Wizard. To learn more about alarms and thresholds, see Creating CloudWatch Alarms.

### To set an upper threshold alarm for a gateway-VTL's upload buffer

1. Start the Create Alarm Wizard.

   a. In the Amazon CloudWatch console, click the **Alarms** link in the **Navigation** pane.
   b. In the **Your CloudWatch Alarms** pane, click **Create Alarm**.

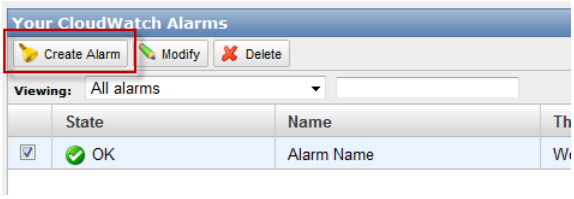


2. Specify a metric for your alarm.

   a. In the **SELECT METRIC** page of the Create Alarm Wizard, select the **AWS/StorageGateway:GatewayId,GatewayName** dimension and find the gateway-VTL that you want to work with.
   b. Select the `UploadBufferPercentUsed` metric. Use the `Average` statistic and a period of 5 minutes.

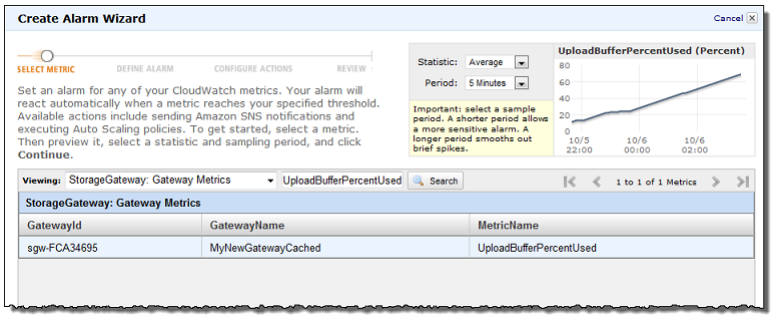
Create Alarm Wizard
Cancel
SELECT METRIC
DEFINE ALARM
CONFIGURE ACTIONS
REVIEW
Set an alarm for any of your CloudWatch metrics. Your alarm will react automatically when a metric reaches your specified threshold. Available actions include sending Amazon SNS notifications and executing Auto Scaling policies. To get started, select a metric. Then preview it, select a statistic and sampling period, and click Continue.
Statistic: Average
Period: 5 Minutes
Important: select a sample period. A shorter period allows a more sensitive alarm. A longer period smooths out brief spikes.
UploadBufferPercentUsed (Percent)
80
60
40
20
0
10/5 22:00
10/6 00:00
10/6 02:00
Viewing: StorageGateway: Gateway Metrics
UploadBufferPercentUsed
Search
1 to 1 of 1 Metrics
StorageGateway: Gateway Metrics
GatewayId
GatewayName
MetricName
sgw-FCA34695
MyNewGatewayCached
UploadBufferPercentUsed

c. Click **Continue**.

3. Define the alarm name, description, and threshold.

   a. In the **DEFINE ALARM** page of the **Create Alarm Wizard**, identify your alarm by giving it a name and description in the **Name** and **Description** boxes.

   b. Define the alarm threshold.

      The following screenshot shows an alarm state that is defined for `UploadBufferPercentUsed` as greater than or equal to 50 percent for 5 minutes.

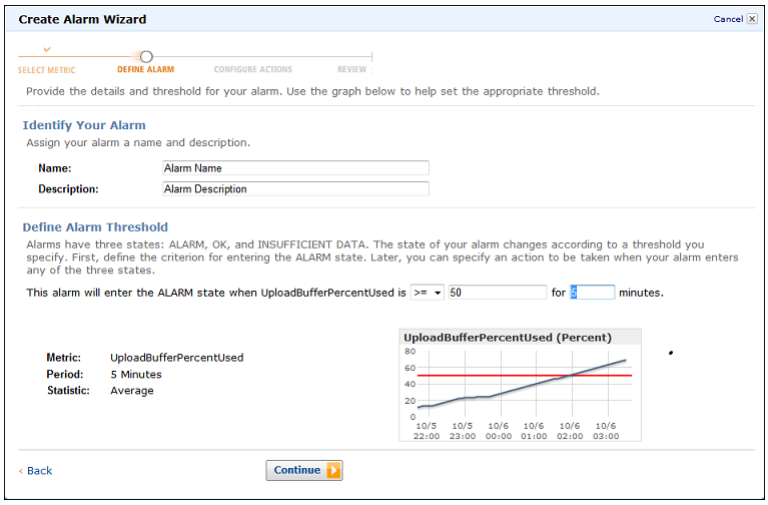
Create Alarm Wizard
Cancel
SELECT METRIC
DEFINE ALARM
CONFIGURE ACTIONS
REVIEW
Provide the details and threshold for your alarm. Use the graph below to help set the appropriate threshold.
Identify Your Alarm
Assign your alarm a name and description.
Name: Alarm Name
Description: Alarm Description
Define Alarm Threshold
Alarms have three states: ALARM, OK, and INSUFFICIENT DATA. The state of your alarm changes according to a threshold you specify. First, define the criterion for entering the ALARM state. Later, you can specify an action to be taken when your alarm enters any of the three states.
This alarm will enter the ALARM state when UploadBufferPercentUsed is >= 50 for 5 minutes.
Metric: UploadBufferPercentUsed
Period: 5 Minutes
Statistic: Average
UploadBufferPercentUsed (Percent)
80
60
40
20
0
10/5 22:00
10/5 23:00
10/6 00:00
10/6 01:00
10/6 02:00
10/6 03:00
Back
Continue

   c. Click **Continue**.

4. Configure an email action for the alarm.

   a. In the **CONFIGURE ACTIONS** page of the **Create Alarm Wizard**, select **ALARM** from the **Alarm State** drop-down list.

   b. Select **Select or create email topic** from the **Topic** drop-down list.

      To create an email topic means that you set up an Amazon Simple Notification Service (Amazon SNS) topic. For more information about Amazon SNS, see Set Up Amazon SNS.

   c. In the **Topic** box, type a descriptive name for the topic.

   d. Click **ADD ACTION**.

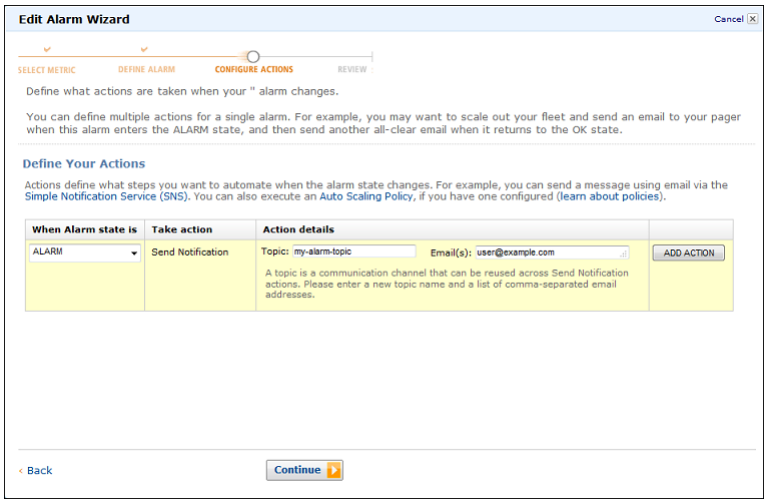
Edit Alarm Wizard
Cancel
SELECT METRIC
DEFINE ALARM
CONFIGURE ACTIONS
REVIEW
Define what actions are taken when your " alarm changes.
You can define multiple actions for a single alarm. For example, you may want to scale out your fleet and send an email to your pager when this alarm enters the ALARM state, and then send another all-clear email when it returns to the OK state.
Define Your Actions
Actions define what steps you want to automate when the alarm state changes. For example, you can send a message using email via the Simple Notification Service (SNS). You can also execute an Auto Scaling Policy, if you have one configured (learn about policies).
When Alarm state is
Take action
Action details
ALARM
Send Notification
Topic: my-alarm-topic
Email(s): user@example.com
ADD ACTION
A topic is a communication channel that can be reused across Send Notification actions. Please enter a new topic name and a list of comma-separated email addresses.
Back
Continue

e. Click **Continue**.

5. Review the alarm settings and create the alarm.

   a. In the **REVIEW** page of the **Create Alarm Wizard**, review the alarm definition, metric, and associated actions for this step.

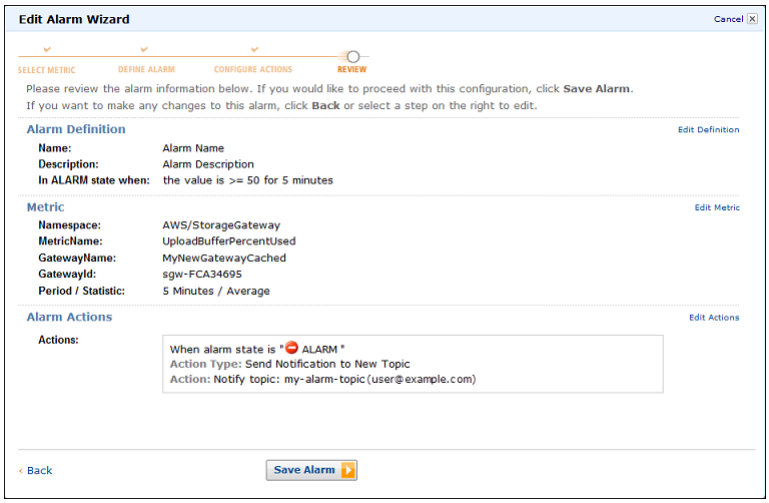
Edit Alarm Wizard
Cancel
SELECT METRIC
DEFINE ALARM
CONFIGURE ACTIONS
REVIEW
Please review the alarm information below. If you would like to proceed with this configuration, click Save Alarm.
If you want to make any changes to this alarm, click Back or select a step on the right to edit.
Alarm Definition
Edit Definition
Name: Alarm Name
Description: Alarm Description
In ALARM state when: the value is >= 50 for 5 minutes
Metric
Edit Metric
Namespace: AWS/StorageGateway
MetricName: UploadBufferPercentUsed
GatewayName: MyNewGatewayCached
GatewayId: sgw-FCA34695
Period / Statistic: 5 Minutes / Average
Alarm Actions
Edit Actions
Actions:
When alarm state is " ALARM "
Action Type: Send Notification to New Topic
Action: Notify topic: my-alarm-topic (user@example.com)
Back
Save Alarm

   b. After reviewing the alarm summary, click **Save Alarm**.

6. Confirm your subscription to the alarm topic.

   a. Open the Amazon Simple Notification Service (Amazon SNS) email to be sent to the email address that you specified when creating the topic.

      The following image shows a notification.

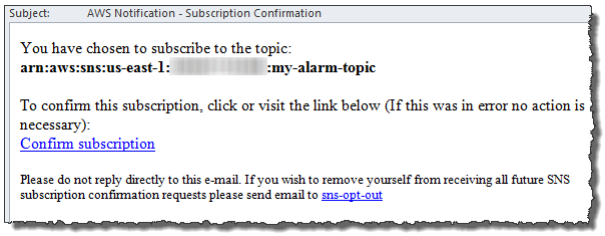
Subject: AWS Notification - Subscription Confirmation
You have chosen to subscribe to the topic:
arn:aws:sns:us-east-1: :my-alarm-topic
To confirm this subscription, click or visit the link below (If this was in error no action is necessary):
Confirm subscription
Please do not reply directly to this e-mail. If you wish to remove yourself from receiving all future SNS subscription confirmation requests please send email to sns-opt-out

b. Confirm your subscription by clicking the link in the email.

   A subscription confirmation displays.

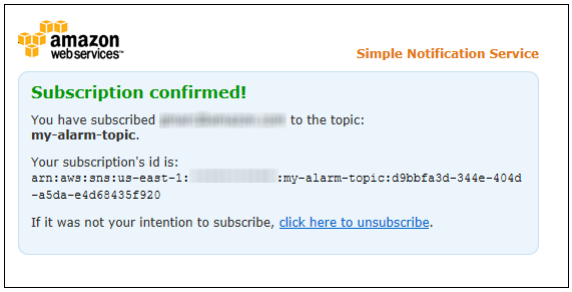
amazon web services
Simple Notification Service
Subscription confirmed!
You have subscribed to the topic:
my-alarm-topic.
Your subscription's id is:
arn:aws:sns:us-east-1: :my-alarm-topic:d9bbfa3d-344e-404d-a5da-e4d68435f920
If it was not your intention to subscribe, click here to unsubscribe.

# Monitoring Cache Storage

The following section discusses how to monitor a gateway-VTL's cache storage and how to create an alarm to get notification when parameters of cache storage exceed specified thresholds. This approach lets you proactively add cache storage to a gateway-VTL.

| Item of Interest | How to Measure |
|---|---|
| Total usage of cache | Use the `CachePercentUsed` and `TotalCacheSize` metrics with the `Average` statistic. For example, use the `CachePercentUsed` with the `Average` statistic to analyze the cache usage over a time.<br><br>The `TotalCacheSize` metric changes only when you add cache to the gateway-VTL. |
| Percentage of read requests that are served from the cache | Use the `CacheHitPercent` metric with the `Average` statistic.<br><br>Typically, you will want `CacheHitPercent` to remain high. If it is not high, and you are reusing old tapes and overwriting stale data, we recommend you delete the old tapes and create new ones. |
| Percentage of cache that is dirty—that is, it contains content that has not been uploaded to AWS | Use the `CachePercentDirty` metrics with the `Average` statistic.<br><br>Typically, you want `CachePercentDirty` to remain low. If it is not low, we recommend adding cache or increasing hardware resources. For more information, see Optimizing Gateway-VTL Performance (p. 341). |

The following procedures assume that you start in the Amazon CloudWatch console.

### To measure the cache's percentage dirty for a gateway-VTL and all its virtual tapes

1. Select the **StorageGateway: Gateway Metrics** dimension, and find the gateway-VTL that you want to work with.
2. Select the `CachePercentDirty` metric.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percentage of cache percentage dirty that is used.

### To measure the cache's percentage dirty for a virtual tape

1. Select the **StorageGateway: TapeMetrics** dimension, and find the gateway-VTL that you want to work with.
2. Select the `CachePercentDirty` metric.
3. Select a **Time Range** value.
4. Select the `Average` statistic.
5. Select a **Period** value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percentage of cache percentage dirty that is used by the tape.

# Understanding Metrics

## Gateway Metrics

In this section, we define *gateway* metrics as metrics that are scoped to the gateway-VTL—that is, they measure something about the gateway-VTL. Because a gateway-VTL contains one or more virtual tapes, a gateway-specific metric is representative of all virtual tapes on the gateway-VTL. For example, the `CloudBytesUploaded` metric is the total number of bytes that the gateway-VTL sent to the cloud during the reporting period. This metric includes the activity of all the tapes on the gateway-VTL.

When working with gateway-VTL metric data, you will specify the unique identification of the gateway-VTL that you are interested in viewing metrics for. To do this, you specify both the `GatewayId` and the `GatewayName` values. When you want to work with metric for a gateway-VTL, you specify the gateway-VTL *dimension* in the metrics namespace, which distinguishes a gateway-specific metric from a tape-specific metric. For more information, see Using Amazon CloudWatch Metrics (p. 352).

The following table describes the AWS Storage Gateway metrics that you can use to get information about your gateway-VTL. The entries in the table are grouped functionally by measure.

**Note**
The reporting period for these metrics is 5 minutes.

| Metric | Description |
|---|---|
| `CacheHitPercent` | Percent of application read operations served from the cache. The sample is taken at the end of the reporting period.<br><br>Units: Percent |

| Metric | Description |
| --- | --- |
| CachePercentUsed | Percent use of the gateway-VTL's cache storage. The sample is taken at the end of the reporting period.<br><br>Units: Percent |
| CachePercentDirty | Percent of the gateway-VTL's cache that has not been persisted to AWS. The sample is taken at the end of the reporting period.<br><br>Units: Percent |
| CloudBytesDownloaded | The total number of precompressed bytes that the gateway-VTL downloaded from AWS during the reporting period.<br><br>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure input/output operations per second (IOPS).<br><br>Units: Bytes |
| CloudDownloadLatency | The total number of milliseconds spent reading data from AWS during the reporting period.<br><br>Use this metric with the Average statistic to measure latency.<br><br>Units: Milliseconds |
| CloudBytesUploaded | The total number of precompressed bytes that the gateway-VTL uploaded to AWS during the reporting period.<br><br>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.<br><br>Units: Bytes |
| UploadBufferFree | The total amount of unused space in the gateway-VTL's upload buffer. The sample is taken at the end of the reporting period.<br><br>Units: Bytes |
| UploadBufferPercentUsed | Percent use of the gateway-VTL's upload buffer. The sample is taken at the end of the reporting period.<br><br>Units: Percent |
| UploadBufferUsed | The total number of bytes being used in the gateway-VTL's upload buffer. The sample is taken at the end of the reporting period.<br><br>Units: Bytes |
| QueuedWrites | The number of bytes waiting to be written to AWS, sampled at the end of the reporting period for all tapes in the gateway-VTL. These bytes are kept in your gateway-VTL's working storage.<br><br>Units: Bytes |

| Metric | Description |
|---|---|
| `ReadBytes` | The total number of bytes read from your gateway-VTL in the reporting period for all tapes in the gateway-VTL.<br><br>Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS.<br><br>Units: Bytes |
| `ReadTime` | The total number of milliseconds spent to do read operations from your gateway-VTL in the reporting period for all tapes in the gateway-VTL.<br><br>Use this metric with the `Average` statistic to measure latency.<br><br>Units: Milliseconds |
| `TotalCacheSize` | The total size of the cache in bytes. The sample is taken at the end of the reporting period.<br><br>Units: Bytes |
| `WriteBytes` | The total number of bytes written to your gateway-VTL in the reporting period for all tapes in the gateway-VTL.<br><br>Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS.<br><br>Units: Bytes |
| `WriteTime` | The total number of milliseconds spent to do write operations to your gateway-VTL in the reporting period for all tapes in the gateway-VTL.<br><br>Use this metric with the `Average` statistic to measure latency.<br><br>Units: Milliseconds |

# Logging Gateway-VTL API Calls by Using AWS CloudTrail

AWS Storage Gateway is integrated with AWS CloudTrail, a service that captures API calls made by or on behalf of AWS Storage Gateway in your AWS account and delivers the log files to an Amazon S3 bucket that you specify. CloudTrail captures API calls from the AWS Storage Gateway console or from the AWS Storage Gateway API. Using the information collected by CloudTrail, you can determine what request was made to AWS Storage Gateway, the source IP address from which the request was made, who made the request, when it was made, and so on. To learn more about CloudTrail, including how to configure and enable it, see the *AWS CloudTrail User Guide*.

## AWS Storage Gateway Information in CloudTrail

When CloudTrail logging is enabled in your AWS account, API calls made to AWS Storage Gateway actions are tracked in log files. AWS Storage Gateway records are written together with other AWS service records in a log file. CloudTrail determines when to create and write to a new file based on a time period and file size.

All of the AWS Storage Gateway actions are logged and are documented in the Actions topic. For example, calls to the ActivateGateway, ListGateways, and ShutdownGateway actions generate entries in the CloudTrail log files.

Every log entry contains information about who generated the request. The user identity information in the log helps you determine whether the request was made with root or IAM user credentials, with temporary security credentials for a role or federated user, or by another AWS service. For more information, go to the **userIdentity** field in the CloudTrail Event Reference in the *AWS CloudTrail User Guide*.

You can store your log files in your bucket for as long as you want, but you can also define Amazon S3 lifecycle rules to archive or delete log files automatically. By default, your log files are encrypted by using Amazon S3 server-side encryption (SSE).

You can choose to have CloudTrail publish Amazon Simple Notification Service (Amazon SNS) notifications when new log files are delivered if you want to take quick action upon log file delivery. For more information, see Configuring Amazon SNS Notifications.

You can also aggregate AWS Storage Gateway log files from multiple AWS regions and multiple AWS accounts into a single Amazon S3 bucket. For more information, see Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket.

## Understanding AWS Storage Gateway Log File Entries

CloudTrail log files can contain one or more log entries where each entry is made up of multiple JSON-formatted events. A log entry represents a single request from any source and includes information about the requested action, any parameters, the date and time of the action, and so on. The log entries are not guaranteed to be in any particular order. That is, they are not an ordered stack trace of the public API calls.

The following example shows a CloudTrail log entry that demonstrates the ActivateGateway action.

```
      {"Records":[{"eventVersion":"1.02","userIdentity":{"type":"IAMUser","prin
cipalId":"AIDAII5AUEPBH2M7JTNVC","arn":"arn:aws:iam::111122223333:user/StorageG
ateway-team/JohnDoe",
       "accountId":"111122223333","accessKeyId":"AKIAIOSFODNN7EXAMPLE","user
Name":"JohnDoe"},"eventTime":"2014-12-04T16:19:00Z",
      "eventSource":"storagegateway.amazonaws.com","eventName":"ActivateGate
way","awsRegion":"us-east-1","sourceIPAddress":"192.0.2.0",
       "userAgent":"aws-cli/1.6.2 Python/2.7.6 Linux/2.6.18-164.el5",
       "requestParameters":{"gatewayTimezone":"GMT-5:00","gateway
Name":"cloudtrailgatewayvtl","gatewayRegion":"us-east-1","activationKey":"EHFBX-
1NDD0-P0IVU-PI259-DHK88","gatewayType":"VTL"},
       "responseElements":{"gatewayARN":"arn:aws:storagegateway:us-east-
1:111122223333:gateway/cloudtrailgatewayvtl"},"requestID":"54BT
FGNQI71987UJD2IHTCT8NF1Q8GLLE1QEU3KPGG6F0KSTAUU0",
       "eventID":"635f2ea2-7e42-45f0-bed1-8b17d7b74265","eventType":"AwsApiC
all","apiVersion":"20130630","recipientAccountId":"444455556666"}]}
```

The following example shows a CloudTrail log entry that demonstrates the ListGateways action.

```
      {"Records":[{"eventVersion":"1.02","userIdentity":{"type":"IAMUser","prin
```

```
cipalId":"AIDAII5AUEPBH2M7JTNVC","arn":"arn:aws:iam::111122223333:user/StorageG
ateway-team/JohnDoe",
        "accountId:"111122223333", "accessKeyId":"AKIAIOSFODNN7EXAMPLE","user
Name":"JohnDoe"},"eventTime":"2014-12-03T19:41:53Z",
        "eventSource":"storagegateway.amazonaws.com","eventName":"ListGate
ways","awsRegion":"us-east-1","sourceIPAddress":"192.0.2.0",
        "userAgent":"aws-cli/1.6.2 Python/2.7.6 Linux/2.6.18-164.el5",
        "requestParameters":null,"responseElements":null,"request
ID":"6U2N42CU37KAO8BG6V1I23FRSJ1Q8GLLE1QEU3KPGG6F0KSTAUU0",
        "eventID":"f76e5919-9362-48ff-a7c4-d203a189ec8d","eventType":"AwsApiC
all","apiVersion":"20130630",
        "recipientAccountId":"444455556666"}]}
```


# Troubleshooting Your Gateway-VTL

This section discusses how to troubleshoot issues related to gateway-VTLs and virtual tapes. The approaches discussed in the section about troubleshooting virtual tapes apply to all gateway-VTL types.

**Topics**


- Troubleshooting Issues for a Gateway-VTL Deployed On-Premises (p. 363)
- Troubleshooting Issues for a Gateway-VTL Deployed on an Amazon EC2 Instance (p. 367)
- Troubleshooting Virtual Tape Issues (p. 370)


## Troubleshooting Issues for a Gateway-VTL Deployed On-Premises

The following table lists typical issues that you might encounter when working with your on-premises gateway-VTLs.

| Issue | Action to Take |
| --- | --- |
| You cannot find the IP address of your gateway-VTL. | Use the hypervisor client to connect to your host to find the gateway-VTL IP address.<br>• For VMware ESXi, the VM's IP address can be found in the vSphere client on the **Summary** tab. Step 2.4: Activate the AWS Storage Gateway (p. 30)<br>• For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console. (see Step 2.4: Activate Your Storage Gateway (p. 50))<br><br>If you are still having trouble finding the gateway-VTL IP address:<br>• Check that the VM is turned on. Only when the VM is turned on does an IP address get assigned to your gateway-VTL.<br>• Wait for the VM to finish starting up. If you just turned on your VM, then it might take several minutes for the gateway-VTL to finish its boot sequence. |

| Issue | Action to Take |
| --- | --- |
| Your gateway-VTL's activation fails when you click the **Proceed to Activation** button in the AWS Storage Gateway console. | • Check that the gateway-VTL VM can be accessed by pinging the VM from your client.<br>• Check that your VM has network connectivity to the Internet. If not, you'll need to configure a Socket Secure (SOCKS) proxy. For more information, see Routing Your Gateway-VTL Through a Proxy (p. 325).<br>• Check that the host has the correct time and is configured to synchronize its time automatically to a Network Time Protocol (NTP) server and that the gateway-VTL VM has the correct time. For information about synchronizing the time of hypervisor hosts and VMs, see Synchronizing Your Gateway-VTL's VM Time (p. 330).<br>• After performing these steps, you can retry the gateway-VTL deployment using the AWS Storage Gateway console and the **Setup and Activate Gateway** wizard.<br>• Check that your VM has at least 7.5 GB of RAM. Gateway allocation fails if there is less than 7.5 GB of RAM. For more information, see Requirements (p. 11). |
| You need to remove a disk allocated as upload buffer space because you want to reduce the amount of upload buffer space for a gateway-VTL. Or, you need to replace a disk used as an upload buffer that has failed. | For instructions about removing a disk allocated as upload buffer space, see Removing Upload Buffer Capacity (Gateway-VTL) (p. 307). |
| You need to improve bandwidth between your gateway-VTL and AWS. | You can improve the bandwidth from your gateway-VTL to AWS by setting up your Internet connection to AWS on a network adapter separate from that of the connection between your applications and the gateway-VTL VM. Doing this is useful if you have a high-bandwidth connection to AWS and you want to avoid bandwidth contention, especially during a snapshot restore. For high-throughput workload needs, you can use AWS Direct Connect to establish a dedicated network connection between your on-premises gateway-VTL and AWS.<br><br>To measure the bandwidth of the connection from your gateway-VTL to AWS, use the `CloudBytesDownloaded` and `CloudBytesUploaded` metrics of the gateway-VTL. For more information, see Measuring Performance Between Your Gateway-VTL and AWS (p. 353). Improving your Internet connectivity helps to ensure that your upload buffer does not fill up. |

| Issue | Action to Take |
| --- | --- |
| Throughput to or from your gateway-VTL drops to zero. | • On the **Gateway** tab of the AWS Storage Gateway console, verify that the IP addresses for your gateway-VTL VM are the same that you see using your hypervisor client software (that is, VMware Vsphere client or Microsoft Hyper-V Manager). If you find a mismatch, restart your gateway-VTL from the AWS Storage Gateway console as shown in Shutting Down and Starting Your Gateway-VTL (p. 311). After the restart, the **IP Addresses** box in the **Gateway** tab of the AWS Storage Gateway console should match the IP addresses for your gateway-VTL that you see in the hypervisor client.<br>  • For VMware ESXi, the VM's IP address can be found in the vSphere client on the **Summary** tab. Step 2.4: Activate the AWS Storage Gateway (p. 30)<br>  • For Microsoft Hyper-V, the VM's IP address can be found by logging in to the local console. Step 2.4: Activate Your Storage Gateway (p. 50)<br>• Check your gateway-VTL's connectivity to AWS, as described in Testing Your Gateway-VTL's Connectivity to the Internet (p. 329).<br>• Check your gateway-VTL's network adapter configuration, and ensure that all the interfaces you intended to be enabled for the gateway-VTL are enabled. To view the network adapter configuration for your gateway-VTL, follow the instructions for Configuring Your Gateway-VTL to Use a Static IP Address (p. 326) and select the option to view your gateway-VTL's network configuration.<br><br>You can view the throughput to and from your gateway-VTL from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway-VTL to AWS, see Measuring Performance Between Your Gateway-VTL and AWS (p. 353). |
| You are having trouble importing (deploying) AWS Storage Gateway on Microsoft Hyper-V. | See Troubleshooting Your Microsoft Hyper-V Setup (p. 417), which discusses some of the common issues you will run into when deploying a gateway-VTL on Microsoft Hyper-V. |

## Allowing Your Gateway-VTL Access Through Firewalls and Routers

Your locally deployed gateway-VTL requires access to the following endpoints to communicate with AWS. If you use firewalls or routers to filter or limit network traffic, you must configure your firewall and router to allow these service endpoints for outbound communication to AWS.

```
client-cp.storagegateway.region.amazonaws.com:443
dp-1.storagegateway.region.amazonaws.com:443
anon-cp.storagegateway.region.amazonaws.com:443
proxy-app.storagegateway.region.amazonaws.com:443
storagegateway.region.amazonaws.com:443
```

The following table provides a list of region strings for the regions available for AWS Storage Gateway–VTL.

| Region Name | Region String |
| --- | --- |
| US East (N. Virginia) | `us-east-1` |

| Region Name | Region String |
|---|---|
| US West (N. California) | `us-west-1` |
| US West (Oregon) | `us-west-2` |
| EU (Ireland) | `eu-west-1` |
| EU (Frankfurt) | `eu-central-1` |
| Asia Pacific (Sydney) | `ap-southeast-2` |
| Asia Pacific (Tokyo) | `ap-northeast-1` |

Depending on your gateway-VTL's region, you replace *`region`* in the endpoint with the corresponding region string. For example, if you create a gateway-VTL in the US West (Oregon) region, the endpoint looks like this: `storagegateway.us-west-2.amazonaws.com:443`.

## Enabling AWS Support Access to Your Gateway-VTL Hosted On-Premises

AWS Storage Gateway provides a local console you can use to perform several maintenance tasks, including enabling AWS Support to access your gateway-VTL to assist you with troubleshooting gateway-VTL issues. By default, AWS Support's access to your gateway-VTL is disabled. You enable this access through the host's local console. To give AWS Support access to your gateway-VTL, you first log in to the local console for the host, navigate to the storage gateway's console, and then connect to the support server.

### To enable AWS Support access to your gateway-VTL

1. Log in to your host's local console. For instructions, go to Logging Into Your Gateway Local Console.

   The local console looks like the following.

   

2. At the prompt, type **5** to open the AWS Storage Gateway console.

   

3. Type **h** to open the **AVAILABLE COMMANDS** window.
4. In the **AVAILABLE COMMANDS** window, type **`open-support-channel`** to connect to customer support for AWS Storage Gateway.
5. At the prompt, type a support service number for your gateway-VTL to connect to the AWS Storage Gateway support server. Choose any number between 2000 and 65000.

Once the support channel is established, provide the support service number you chose to AWS Support so AWS Support can provide troubleshooting assistance.

6. When the support session is completed, type **q** to end it.
7. Type **exit** to log out of the AWS Storage Gateway console.
8. Follow the prompts to exit the local console.

## Troubleshooting Issues for a Gateway-VTL Deployed on an Amazon EC2 Instance

The following table lists typical issues that you might encounter working with your gateway-VTL deployed on Amazon Elastic Compute Cloud (Amazon EC2). For more information about the difference between an on-premises gateway-VTL and a gateway-VTL deployed in Amazon EC2, see Amazon EC2 Gateway (p. 57).

| Issue | Action to Take |
|---|---|
| Your Amazon EC2 gateway-VTL activation fails when you click the **Proceed to Activation** button in the AWS Storage Gateway console. | If activation has not occurred in a few moments, check the following in the Amazon EC2 console:<br>• Port 80 is enabled in the security group you associated with the instance. For more information about adding a security group rule, go to Adding a Security Group Rule in the *Amazon EC2 User Guide for Linux Instances*.<br>• The gateway-VTL instance is marked as running. In the Amazon EC2 console, the **State** of the instance should be "running".<br><br>After correcting the problem, try activating the gateway-VTL again by going to the AWS Storage Gateway console, clicking **Deploy a new Gateway on Amazon EC2**, and re-entering the IP address of the instance. |

| Issue | Action to Take |
| --- | --- |
| You can't find your Amazon EC2 gateway-VTL instance in the list of instances. | In this case, you can take the following actions to find the gateway-VTL instance:<br>• Check the name of the Amazon Machine Image (AMI) name in the **Description** tab of the instance. An instance based on the AWS Storage Gateway AMI should start with the text **aws-storage-gateway-ami**.<br>• If you have several instances that are based on the AWS Storage Gateway AMI, check the instance launch time to find the correct instance.<br><br>You might not be able to find your instance if you did not give your instance a resource tag and you have many instances running. In this situation, it can be hard to tell which instance you deployed the gateway-VTL in. |
| You created an Amazon EBS volume but can't attach it to your Amazon EC2 gateway-VTL instance. | Check that the Amazon EBS volume in question is in the same Availability Zone as the gateway-VTL instance. If there is a discrepancy in Availability Zones, create a new Amazon EBS volume in the same Availability Zone as your instance. |
| You can't attach an initiator to a storage volume target of your Amazon EC2 gateway-VTL. | Check that the security group you launched the instance with includes a rule allowing the port you are using for iSCSI. The port is usually set as 3260. For more information about connecting to storage volumes, see Connecting to Storage Volumes (p. 111). |
| You activated your Amazon EC2 gateway-VTL, but when you attempt to add storage volumes, you receive an error message indicating you have no disks available. | For a newly activated gateway-VTL, no volume storage is defined. Before you can define volume storage, you must allocate local disks to the gateway-VTL to use as upload buffer space and cache storage.<br><br>For a gateway-VTL deployed to Amazon EC2, the local disks are Amazon EBS volumes attached to the instance. This error message likely occurs because no Amazon EBS volumes are defined for the instance. Check block devices defined for the instance that is running the gateway-VTL. If there are only two block devices (the default devices that come with the AMI), then you should add storage (see Amazon EC2 Gateway (p. 57)). After attaching two or more Amazon EBS volumes, try creating volume storage on the gateway-VTL. |
| You need to remove a disk allocated as upload buffer space, because you want to reduce the amount of upload buffer space. | Follow the steps in To remove an Amazon EBS volume from your Amazon EC2–hosted gateway (p. 161). |

| Issue | Action to Take |
| --- | --- |
| Throughput to or from your Amazon EC2 gateway-VTL drops to zero. | • Verify the gateway-VTL instance is running. For example, if the instance is starting due to a reboot, wait for the instance to restart.<br>• Verify that the gateway-VTL IP address has not changed. If the instance was stopped and then restarted, the IP address of the instance might have changed. In this case, you need to activate a new gateway-VTL.<br><br>You can view the throughput to and from your gateway-VTL from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway-VTL to AWS, see Measuring Performance Between Your Gateway and AWS (p. 188). |

**Warning**
For gateway-VTLs that are deployed on an Amazon EC2 instance, accessing the gateway-VTL over a public Internet connection is not supported. The elastic IP address of the Amazon EC2 instance cannot be used as the target address.

## Enabling AWS Support to Access a Gateway Hosted on an Amazon EC2 Instance

AWS Storage Gateway provides a local console you can use to perform several maintenance tasks, including enabling AWS Support to access your gateway-VTL to assist you with troubleshooting gateway-VTL issues. By default, AWS Support's access to your gateway-VTL is disabled. You enable this access through the Amazon EC2 local console. You log in to the Amazon EC2 local console via Secure Shell (SSH). To successfully log in via SSH, your instance's security group must have a rule allowing access via TCP port 22.

**Note**
If you add a new rule to an existing security group, the new rule applies to all instances that use that security group. For more information about security groups and how to add a security group rule, go to Amazon EC2 Security Groups in the *Amazon EC2 User Guide.*

To let AWS Support connect to your gateway-VTL, you first log in to the local console for the Amazon EC2 instance, navigate to the storage gateway-VTL's console, and then provide the access.

### To enable AWS support access to a gateway-VTL deployed on an Amazon EC2 instance

1. Log in to the local console for your Amazon EC2 instance. For instructions, go to Connect to Your Instance.

   You can use the following command to log in to the EC2 instance's local console.

```
ssh –i PRIVATE-KEY sguser@INSTANCE-PUBLIC-DNS-NAME
```

   **Note**
   The *PRIVATE-KEY* is the `.pem` file containing the private certificate of the EC2 key pair that you used to launch the Amazon EC2 instance. For more information, go to Retrieving the Public Key for Your Key Pair in the *Amazon EC2 User Guide*.
   The *INSTANCE-PUBLIC-DNS-NAME* is the public Domain Name System (DNS) name of your Amazon EC2 instance that your gateway-VTL is running on. You obtain this public DNS name by selecting the Amazon EC2 instance in the EC2 console and clicking the **Description** tab.

The local console looks like the following.

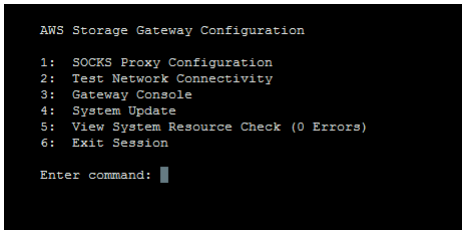


2. At the prompt, type **3** to open the AWS Storage Gateway console.

3. Type **h** to open the **AVAILABLE COMMANDS** window.
4. In the **AVAILABLE COMMANDS** window, type **`open-support-channel`** to connect to customer support for AWS Storage Gateway.
5. At the prompt, type a support service number for your gateway-VTL to connect to the AWS Storage Gateway support server. Choose any number between 2000 and 65000.

Once the support channel is established, provide the support service number you selected to AWS Support so they can provide troubleshooting assistance.

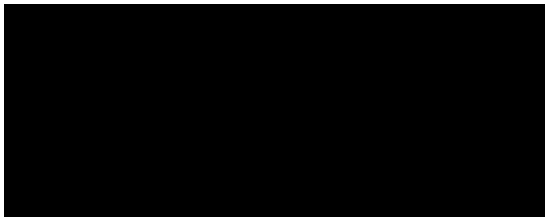

6. When the support session is completed, type **q** to end it.
7. Type **`exit`** to exit the AWS Storage Gateway console.
8. Follow the console menus to log out of the AWS Storage Gateway instance.

## Troubleshooting Virtual Tape Issues

This section discusses actions to take if you experience unexpected issues with your virtual tapes.

## Irrecoverable Tapes

If your virtual tape fails unexpectedly, AWS Storage Gateway sets the status of the failed virtual tape to IRRECOVERABLE. The action you take depends on the circumstance.

The following table list possible circumstances and the actions to take.

| Issue | Action to Take |
| --- | --- |
| A virtual tape has the status IRRECOVERABLE, and you don't need it. This tape has never been archived to VTS. | Delete the tape. For more information, see Deleting Tapes from Your Gateway-VTL (p. 280). |
| A virtual tape has status IRRECOVERABLE, and you don't need it. This tape has previously been archived to VTS and retrieved to your gateway-VTL. | Use your backup software to archive the tape in VTS. For more information, see Archiving Tapes to the VTS (p. 281). |
| A virtual tape has status IRRECOVERABLE, and you need to work with it. | 1. Activate a new gateway-VTL if you don't have one activated. For more information, see Step 2: Deploy a VM and Activate the Gateway (p. 14).<br>2. Disable the gateway-VTL that contains the irrecoverable tape, and recover the tape from a recovery point to the new gateway-VTL. For more information, see Recovering a Virtual Tape from a Malfunctioning Gateway-VTL (p. 282).<br>**Note**<br>You have to reconfigure your iSCSI initiator and backup application to use the new gateway-VTL. For more information, see Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245). |

| Issue | Action to Take |
| --- | --- |
| A cache disk in your gateway encounters a failure. | If one or more cache disks in your gateway encounters a failure, the gateway prevents read and write operations to your virtual tapes and volumes. To resume normal functionality, reconfigure your gateway as described following:<br>• If the cache disk is inaccessible or unusable, delete the disk from your gateway configuration.<br>• If the cache disk is still accessible and useable, reconnect it to your gateway.<br>**Note**<br>If you delete a cache disk, tapes or volumes that have clean data (that is, for which data in the cache disk and Amazon S3 are synchronized) will continue to be available when the gateway resumes normal functionality. For example, if your gateway has three cache disks and you delete two, tapes or volumes that are clean will have AVAILABLE status. Other tapes and volumes will have IRRECOVERABLE status.<br>If you use ephemeral disks as cache disks for your gateway or mount your cache disks on an ephemeral drive, your cache disks will be lost when you shut down the gateway. Shutting down the gateway when your cache disk and Amazon S3 are not synchronized can result in data loss. As a result, we don't recommend using ephemeral drives or disks. |

# Access Control Using AWS Identity and Access Management (IAM)

AWS Identity and Access Management (IAM) helps you securely control access to Amazon Web Services and your account resources. With IAM, you can create multiple IAM users under the umbrella of your AWS account. To learn more about IAM and its features, go to What Is IAM?

Every user you create in the IAM system starts with no permissions. In other words, by default, users can do nothing. A *permission* is a general term we use to mean the ability to perform an action against a resource. The AWS Storage Gateway API (see API Reference for AWS Storage Gateway (p. 379)) enables a list of actions you can perform. However, unless you explicitly grant a user permissions, that user cannot perform any of these actions. You grant a permission to a user with a policy. A policy is a document that formally states one or more permissions. For more information about IAM policies, go to Overview of Policies.

You write a policy using the access policy language that IAM uses. You then attach the policy to a user or a group in your AWS account. For more information about the policy language, go to The Access Policy Language in *IAM User Guide*.

The Element Descriptions section of *IAM User Guide* describes elements you can use in a policy. The following information about some of the policy elements is specific to AWS Storage Gateway:

- **Resource—**The object or objects the policy covers. You identify resources using the following Amazon Resource Name (ARN) format.

```
arn:aws:<vendor>:<region>:<namespace>:<relative-id>
```

In this format, *vendor* is the product name "storagegateway" and *namespace* is the account ID. In AWS Storage Gateway, there are three types of resources, `gateway`, `volume`, and `iSCSITarget`. For each type of resource, the following table shows example ARNs.

| Resource | Description |
|---|---|
| Gateway ARN | arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway |
| Volume ARN | arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB |

| Resource | Description |
|---|---|
| Target ARN (name of an iSCSI target) | arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume |

- **Action—**The specific type or types of action allowed or denied. For a complete list of AWS Storage Gateway actions, see Operations in AWS Storage Gateway (p. 403).

  **Note**
  The Amazon EBS snapshots generated from AWS Storage Gateway are managed as Amazon EC2 resources and corresponding EC2 actions.

This section provides example IAM policies that illustrate how to grant a user permission to perform specific AWS Storage Gateway actions. You can then attach these policies to a user for whom you want to grant access permissions.

# Example Policies

### Example 1: Allow all actions

The following policy allows a user to perform all the AWS Storage Gateway actions. The policy also allows the user to perform Amazon EC2 actions (DescribeSnapshots and DeleteSnapshot) on the Amazon EBS snapshots generated from AWS Storage Gateway.

```
{
   "Statement":[
      {
         "Action": [
            "storagegateway:*"
         ],
         "Effect": "Allow",
         "Resource": "*"
      },
      {
         "Action": [
            "ec2:DescribeSnapshots",
            "ec2:DeleteSnapshot"
         ],
         "Effect": "Allow",
         "Resource": "*"
      }
   ]
}
```

### Example 2: Allow read-only access to a gateway

The following policy allows all `List*` and `Describe*` actions on all resources. Note that these actions are read actions. So the policy does not allow the user to change state of any resources—that is, the policy does not allow the user to perform the actions such as `DeleteGateway`, `ActivateGateway`, and `ShutdownGateway`.

The policy also allows the `DescribeSnapshots` Amazon EC2 action. For more information, go to DescribeSnapshots in the *Amazon EC2 API Reference*.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:List*",
        "storagegateway:Describe*"
      ],
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Action": [
        "ec2:DescribeSnapshots"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

In the preceding policy, instead of a using a wild card, you could scope resources covered by the policy to a specific gateway. The policy would then allow the actions only on the specific gateway.

```
"Resource": "arn:aws:storagegateway:us-east-1:111122223333:gateway/[Gateway Name]/*"
```

Within a gateway, you can further restrict the scope of the resources to only the gateway volumes.

```
"Resource": "arn:aws:storagegateway:us-east-1:111122223333:gateway/[Gateway Name]/volume/*"
```

### Example 3: Allow access to a specific gateway

The following policy allows all actions on a specific gateway. That is, the user is restricted from accessing other gateways you might have deployed.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gate
way/[Gateway Name]/*"
    }
  ]
}
```

The preceding policy works if the user to whom the policy is attached uses either the API or an AWS SDK to access the gateway. However, if this user plans to use the AWS Storage Gateway console, you must also grant permission to the `ListGateways` action.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gate
way/[Gateway Name]/*"
    },
    {
      "Action": [
        "storagegateway:ListGateways"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

Additionally, if the user plans to activate the specific gateway, you must also grant permission to the `ActivateGateway` action.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gate
way/[Gateway Name]/*"
    },
    {
```

```
      "Action": [
        "storagegateway:ListGateways",
        "storagegateway:ActivateGateway"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

### Example 4: Grant permissions to access a specific volume

The following policy allows a user all actions to a specific volume on a gateway. Because a user does not get any permissions by default, the policy restricts the user to accessing only a specific volume.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gate
way/[Gateway Name]/volume/[Volume Name]"
    }
  ]
}
```

The preceding policy works if the user to whom the policy is attached uses either the API or an AWS SDK to access the volume. However, if this user plans to use the AWS Storage Gateway console, you must also grant permission to the `ListGateways` action.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gate
way/[Gateway Name]/volume/[Volume Name]"
    },
    {
      "Action": [
        "storagegateway:ListGateways"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

### Example 5: Allow all actions on gateways with a specific prefix

The following policy allows a user to perform all action on gateways whose name starts with "DeptX". The policy also allows the `DescribeSnapshots` Amazon EC2 action.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gate
way/[Gateway Name Prefix]*"
    },
    {
      "Action": [
        "ec2:DescribeSnapshots"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

The preceding policy works if the user to whom the policy is attached uses either the API or an AWS SDK to access the gateway. However, if this user plans to use the AWS Storage Gateway console, you must grant additional permissions as described in 3: Allow access to a specific gateway (p. 376).

# API Reference for AWS Storage Gateway

In addition to using the console, you can use the AWS Storage Gateway API to programmatically configure and manage your gateways. This section describes the AWS Storage Gateway operations, request signing for authentication and the error handling. For information about the regions and endpoints available for AWS Storage Gateway, see Regions and Endpoints.

> **Note**
> You can also use the AWS SDKs when developing applications with AWS Storage Gateway. The AWS SDKs for Java, .NET, and PHP wrap the underlying AWS Storage Gateway API, simplifying your programming tasks. For information about downloading the SDK libraries, go to Sample Code Libraries.

**Topics**


- AWS Storage Gateway Required Request Headers (p. 379)
- Signing Requests (p. 381)
- Error Responses (p. 383)
- Operations in AWS Storage Gateway (p. 403)


## AWS Storage Gateway Required Request Headers

This section describes the required headers that you must send with every POST request to AWS Storage Gateway. You include HTTP headers to identify key information about the request including the operation you want to invoke, the date of the request, and information that indicates the authorization of you as the sender of the request. Headers are case insensitive and the order of the headers is not important.

The following example shows headers that are used in the API_ActivateGateway operation.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-
```

```
east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-
amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066ab
dd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ActivateGateway
```

The following are the headers that must include with your POST requests to AWS Storage Gateway. Headers shown below that begin with "x-amz" are AWS-specific headers. All other headers listed are common header used in HTTP transactions.

| Header | Description |
|---|---|
| `Authorization` | The authorization header contains several of pieces of information about the request that enable AWS Storage Gateway to determine if the request is a valid action for the requester. The format of this header is as follows (line breaks added for readability):<br>`Authorization: AWS4-HMAC_SHA456 Credentials=YourAccessKey/yyymmdd/region/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=CalculatedSignature`<br>In the preceding syntax, you specify *YourAccessKey*, the year, month, and day (*yyyymmdd*), the *region*, and the *CalculatedSignature*. The format of the authorization header is dictated by the requirements of the AWS V4 Signing process. The details of signing are discussed in the topic Signing Requests (p. 381). |
| `Content-Type` | Use `application/x-amz-json-1.1` as the content type for all requests to AWS Storage Gateway.<br>`Content-Type: application/x-amz-json-1.1` |
| `Host` | Use the host header to specify the AWS Storage Gateway endpoint where you send your request. For example, `storagegateway.us-east-1.amazonaws.com` is the endpoint for the US East (N. Virginia) region. For more information about the endpoints available for AWS Storage Gateway, see Regions and Endpoints.<br>`Host: storagegateway.region.amazonaws.com` |
| `x-amz-date` | You must provide the time stamp in either the HTTP `Date` header or the AWS `x-amz-date` header. (Some HTTP client libraries don't let you set the `Date` header.) When an `x-amz-date` header is present, the AWS Storage Gateway ignores any `Date` header during the request authentication. The `x-amz-date` format must be ISO8601 Basic in the YYYYMMDD'T'HHMMSS'Z' format. If both the `Date` and `x-amz-date` header are used, the format of the Date header does not have to be ISO8601.<br>`x-amz-date: YYYYMMDD'T'HHMMSS'Z'` |

| Header | Description |
| --- | --- |
| `x-amz-target` | This header specifies the version of the API and the operation that you are requesting. The target header values are formed by concatenating the API version with the API name and are in the following format.<br>`x-amz-target: StorageGateway_APIversion.operationName`<br>The *operationName* value (e.g. "ActivateGateway") can be found from the API list, API Reference for AWS Storage Gateway (p. 379). |

# Signing Requests

AWS Storage Gateway requires that you authenticate every request you send by signing the request. To sign a request, you calculate a digital signature using a cryptographic hash function. A cryptographic hash is a function that returns a unique hash value based on the input. The input to the hash function includes the text of your request and your secret access key. The hash function returns a hash value that you include in the request as your signature. The signature is part of the `Authorization` header of your request.

After receiving your request, AWS Storage Gateway recalculates the signature using the same hash function and input that you used to sign the request. If the resulting signature matches the signature in the request, AWS Storage Gateway processes the request. Otherwise, the request is rejected.

AWS Storage Gateway supports authentication using AWS Signature Version 4. The process for calculating a signature can be broken into three tasks:

- Task 1: Create a Canonical Request

  Rearrange your HTTP request into a canonical format. Using a canonical form is necessary because AWS Storage Gateway uses the same canonical form when it recalculates a signature to compare with the one you sent.
- Task 2: Create a String to Sign

  Create a string that you will use as one of the input values to your cryptographic hash function. The string, called the *string to sign*, is a concatenation of the name of the hash algorithm, the request date, a *credential scope* string, and the canonicalized request from the previous task. The *credential scope* string itself is a concatenation of date, region, and service information.
- Task 3: Create a Signature

  Create a signature for your request by using a cryptographic hash function that accepts two input strings: your *string to sign* and a *derived key*. The *derived key* is calculated by starting with your secret access key and using the *credential scope* string to create a series of Hash-based Message Authentication Codes (HMACs).

# Example Signature Calculation

The following example walks you through the details of creating a signature for API_ListGateways. The example could be used as a reference to check your signature calculation method. Other reference calculations are included in the Signature Version 4 Test Suite of the Amazon Web Services Glossary.

The example assumes the following:

- The time stamp of the request is "Mon, 10 Sep 2012 00:00:00" GMT.

- The endpoint is the US East (N. Virginia) region.

The general request syntax (including the JSON body) is:

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
x-amz-Date: 20120910T000000Z
Authorization: SignatureToBeCalculated
Content-type: application/x-amz-json-1.1
x-amz-target: StorageGateway_20120630.ListGateways
{}
```

The canonical form of the request calculated for Task 1: Create a Canonical Request (p. 381) is:

```
POST
/

content-type:application/x-amz-json-1.1
host:storagegateway.us-east-1.amazonaws.com
x-amz-date:20120910T000000Z
x-amz-target:StorageGateway_20120630.ListGateways

content-type;host;x-amz-date;x-amz-target
44136fa355b3678a1146ad16f7e8649e94fb4fc21fe77e8310c060f61caaff8a
```

The last line of the canonical request is the hash of the request body. Also, note the empty third line in the canonical request. This is because there are no query parameters for this API (or any AWS Storage Gateway APIs).

The *string to sign* for Task 2: Create a String to Sign (p. 381) is:

```
AWS4-HMAC-SHA256
20120910T000000Z
20120910/us-east-1/storagegateway/aws4_request
92c0effa6f9224ac752ca179a04cecbede3038b0959666a8160ab452c9e51b3e
```

The first line of the *string to sign* is the algorithm, the second line is the time stamp, the third line is the *credential scope*, and the last line is a hash of the canonical request from Task 1.

For Task 3: Create a Signature (p. 381), the *derived key* can be represented as:

```
derived key = HMAC(HMAC(HMAC(HMAC("AWS4" + YourSecretAccessKey,"20120910"),"us-
east-1"),"storagegateway"),"aws4_request")
```

If the secret access key, wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY, is used, then the calculated signature is:

```
6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

The final step is to construct the `Authorization` header. For the demonstration access key AKIAIOSFODNN7EXAMPLE, the header (with line breaks added for readability) is:

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120910/us-
east-1/storagegateway/aws4_request,
SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

# Error Responses

**Topics**


- Exceptions (p. 383)
- Operation Error Codes (p. 384)
- Error Responses (p. 402)


This section provides reference information about AWS Storage Gateway errors. These errors are represented by an error exception and an operation error code. For example, the error exception `InvalidSignatureException` is returned by any API response if there is a problem with the request signature. However, the operation error code `ActivationKeyInvalid` is returned only for the ActivateGateway API.

Depending on the type of error, AWS Storage Gateway may return only just an exception, or it may return both an exception and an operation error code. Examples of error responses are shown in the Error Responses (p. 402).

## Exceptions

The following table lists AWS Storage Gateway API exceptions. When an AWS Storage Gateway operation returns an error response, the response body contains one of these exceptions. The `InternalServerError` and `InvalidGatewayRequestException` return one of the operation error codes Operation Error Codes (p. 384) message codes that give the specific operation error code.

| Exception | Message | HTTP Status Code |
|---|---|---|
| `IncompleteSignatureException` | The specified signature is incomplete. | 400 Bad Request |
| `InternalFailure` | The request processing has failed due to some unknown error, exception or failure. | 500 Internal Server Error |
| `InternalServerError` | One of the operation error code messages Operation Error Codes (p. 384). | 500 Internal Server Error |
| `InvalidAction` | The requested action or operation is invalid. | 400 Bad Request |
| `InvalidClientTokenId` | The X.509 certificate or AWS Access Key ID provided does not exist in our records. | 403 Forbidden |
| `InvalidGatewayRequestException` | One of the operation error code messages in Operation Error Codes (p. 384). | 400 Bad Request |
| `InvalidSignatureException` | The request signature we calculated does not match the signature you provided. Check your AWS Access Key and signing method. | 400 Bad Request |

| Exception | Message | HTTP Status Code |
|---|---|---|
| MissingAction | The request is missing an action or operation parameter. | 400 Bad Request |
| MissingAuthenticationToken | The request must contain either a valid (registered) AWS Access Key ID or X.509 certificate. | 403 Forbidden |
| RequestExpired | The request is past the expiration date or the request date (either with 15 minute padding), or the request date occurs more than 15 minutes in the future. | 400 Bad Request |
| SerializationException | An error occurred during serialization. Check that your JSON payload is well-formed. | 400 Bad Request |
| ServiceUnavailable | The request has failed due to a temporary failure of the server. | 503 Service Unavailable |
| SubscriptionRequiredException | The AWS Access Key Id needs a subscription for the service. | 400 Bad Request |
| ThrottlingException | Rate exceeded. | 400 Bad Request |
| UnknownOperationException | An unknown operation was specified. Valid operations are listed in Operations in AWS Storage Gateway (p. 403). | 400 Bad Request |
| UnrecognizedClientException | The security token included in the request is invalid. | 400 Bad Request |
| ValidationException | The value of an input parameter is bad or out of range. | 400 Bad Request |

# Operation Error Codes

The following table shows the mapping between AWS Storage Gateway operation error codes and APIs that can return the codes. All operation error codes are returned with one of two general exceptions—InternalServerError and InvalidGatewayRequestException—described in Exceptions (p. 383).

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| ActivationKeyExpired | The specified activation key has expired. | ActivateGateway |
| ActivationKeyInvalid | The specified activation key is invalid. | ActivateGateway |
| ActivationKeyNotFound | The specified activation key was not found. | ActivateGateway |
| BandwidthThrottleScheduleNotFound | The specified bandwidth throttle was not found. | DeleteBandwidthRateLimit |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| `CannotExportSnapshot` | The specified snapshot cannot be exported. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume |
| `InitiatorNotFound` | The specified initiator was not found. | DeleteChapCredentials |
| `DiskAlreadyAllocated` | The specified disk is already allocated. | AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateStorediSCSIVolume |
| `DiskDoesNotExist` | The specified disk does not exist. | AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateStorediSCSIVolume |
| `DiskSizeNotGigAligned` | The specified disk is not gigabyte-aligned. | CreateStorediSCSIVolume |
| `DiskSizeGreaterThan-VolumeMaxSize` | The specified disk size is greater than the maximum volume size. | CreateStorediSCSIVolume |
| `DiskSizeLessThanVolumeSize` | The specified disk size is less than the volume size. | CreateStorediSCSIVolume |
| `DuplicateCertificateInfo` | The specified certificate information is a duplicate. | ActivateGateway |

| Operation Error Code | Message | Operations That Return this Error Code |
| --- | --- | --- |
| `GatewayInternalError` | A gateway internal error occurred. | AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateCachediSCSIVolume<br>CreateSnapshot<br>CreateStorediSCSIVolume<br>CreateSnapshotFromVolumeRecoveryPoint<br>DeleteBandwidthRateLimit<br>DeleteChapCredentials<br>DeleteVolume<br>DescribeBandwidthRateLimit<br>DescribeCache<br>DescribeCachediSCSIVolumes<br>DescribeChapCredentials<br>DescribeGatewayInformation<br>DescribeMaintenanceStartTime<br>DescribeSnapshotSchedule<br>DescribeStorediSCSIVolumes<br>DescribeWorkingStorage<br>ListLocalDisks<br>ListVolumes<br>ListVolumeRecoveryPoints<br>ShutdownGateway<br>StartGateway<br>UpdateBandwidthRateLimit<br>UpdateChapCredentials<br>UpdateMaintenanceStartTime<br>UpdateGatewaySoftwareNow<br>UpdateSnapshotSchedule |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| `GatewayNotConnected` | The specified gateway is not connected. | AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateCachediSCSIVolume<br>CreateSnapshot<br>CreateStorediSCSIVolume<br>CreateSnapshotFromVolumeRecoveryPoint<br>DeleteBandwidthRateLimit<br>DeleteChapCredentials<br>DeleteVolume<br>DescribeBandwidthRateLimit<br>DescribeCache<br>DescribeCachediSCSIVolumes<br>DescribeChapCredentials<br>DescribeGatewayInformation<br>DescribeMaintenanceStartTime<br>DescribeSnapshotSchedule<br>DescribeStorediSCSIVolumes<br>DescribeWorkingStorage<br>ListLocalDisks<br>ListVolumes<br>ListVolumeRecoveryPoints<br>ShutdownGateway<br>StartGateway<br>UpdateBandwidthRateLimit<br>UpdateChapCredentials<br>UpdateMaintenanceStartTime<br>UpdateGatewaySoftwareNow<br>UpdateSnapshotSchedule |

| Operation Error Code | Message | Operations That Return this Error Code |
| --- | --- | --- |
| `GatewayNotFound` | The specified gateway was not found. | |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| | | AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateCachediSCSIVolume<br>CreateSnapshot<br>CreateSnapshotFromVolumeRecov-eryPoint<br>CreateStorediSCSIVolume<br>DeleteBandwidthRateLimit<br>DeleteChapCredentials<br>DeleteGateway<br>DeleteVolume<br>DescribeBandwidthRateLimit<br>DescribeCache<br>DescribeCachediSCSIVolumes<br>DescribeChapCredentials<br>DescribeGatewayInformation<br>DescribeMaintenanceStartTime<br>DescribeSnapshotSchedule<br>DescribeStorediSCSIVolumes<br>DescribeWorkingStorage<br>ListLocalDisks<br>ListVolumes<br>ListVolumeRecoveryPoints<br>ShutdownGateway<br>StartGateway<br>UpdateBandwidthRateLimit<br>UpdateChapCredentials<br>UpdateMaintenanceStartTime<br>UpdateGatewaySoftwareNow |

| Operation Error Code | Message | Operations That Return this Error Code |
| --- | --- | --- |
| | | UpdateSnapshotSchedule |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| `GatewayProxyNetworkConnectionBusy` | The specified gateway proxy network connection is busy. | AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateCachediSCSIVolume<br>CreateSnapshot<br>CreateSnapshotFromVolumeRecoveryPoint<br>CreateStorediSCSIVolume<br>DeleteBandwidthRateLimit<br>DeleteChapCredentials<br>DeleteVolume<br>DescribeBandwidthRateLimit<br>DescribeCache<br>DescribeCachediSCSIVolumes<br>DescribeChapCredentials<br>DescribeGatewayInformation<br>DescribeMaintenanceStartTime<br>DescribeSnapshotSchedule<br>DescribeStorediSCSIVolumes<br>DescribeWorkingStorage<br>ListLocalDisks<br>ListVolumes<br>ListVolumeRecoveryPoints<br>ShutdownGateway<br>StartGateway<br>UpdateBandwidthRateLimit<br>UpdateChapCredentials<br>UpdateMaintenanceStartTime<br>UpdateGatewaySoftwareNow<br>UpdateSnapshotSchedule |

| Operation Error Code | Message | Operations That Return this Error Code |
| --- | --- | --- |
| `InternalError` | An internal error occurred. | |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| | | ActivateGateway<br>AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateCachediSCSIVolume<br>CreateSnapshot<br>CreateSnapshotFromVolumeRecoveryPoint<br>CreateStorediSCSIVolume<br>DeleteBandwidthRateLimit<br>DeleteChapCredentials<br>DeleteGateway<br>DeleteVolume<br>DescribeBandwidthRateLimit<br>DescribeCache<br>DescribeCachediSCSIVolumes<br>DescribeChapCredentials<br>DescribeGatewayInformation<br>DescribeMaintenanceStartTime<br>DescribeSnapshotSchedule<br>DescribeStorediSCSIVolumes<br>DescribeWorkingStorage<br>ListLocalDisks<br>ListGateways<br>ListVolumes<br>ListVolumeRecoveryPoints<br>ShutdownGateway<br>StartGateway<br>UpdateBandwidthRateLimit<br>UpdateChapCredentials |

| Operation Error Code | Message | Operations That Return this Error Code |
| --- | --- | --- |
| | | UpdateMaintenanceStartTime<br>UpdateGatewayInformation<br>UpdateGatewaySoftwareNow<br>UpdateSnapshotSchedule |

| Operation Error Code | Message | Operations That Return this Error Code |
| --- | --- | --- |
| `InvalidParameters` | The specified request contains invalid parameters. | |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| | | ActivateGateway<br>AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateCachediSCSIVolume<br>CreateSnapshot<br>CreateSnapshotFromVolumeRecoveryPoint<br>CreateStorediSCSIVolume<br>DeleteBandwidthRateLimit<br>DeleteChapCredentials<br>DeleteGateway<br>DeleteVolume<br>DescribeBandwidthRateLimit<br>DescribeCache<br>DescribeCachediSCSIVolumes<br>DescribeChapCredentials<br>DescribeGatewayInformation<br>DescribeMaintenanceStartTime<br>DescribeSnapshotSchedule<br>DescribeStorediSCSIVolumes<br>DescribeWorkingStorage<br>ListLocalDisks<br>ListGateways<br>ListVolumes<br>ListVolumeRecoveryPoints<br>ShutdownGateway<br>StartGateway<br>UpdateBandwidthRateLimit<br>UpdateChapCredentials |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| | | UpdateMaintenanceStartTime<br>UpdateGatewayInformation<br>UpdateGatewaySoftwareNow<br>UpdateSnapshotSchedule |
| `LocalStorageLimitExceeded` | The local storage limit was exceeded. | AddCache<br>AddUploadBuffer<br>AddWorkingStorage |
| `LunInvalid` | The specified LUN is invalid. | CreateStorediSCSIVolume |
| `MaximumVolumeCountExceeded` | The maximum volume count was exceeded. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume<br>DescribeCachediSCSIVolumes<br>DescribeStorediSCSIVolumes |
| `NetworkConfigurationChanged` | The gateway network configuration has changed. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume |

| Operation Error Code | Message | Operations That Return this Error Code |
| --- | --- | --- |
| `NotSupported` | The specified operation is not supported. | |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| | | ActivateGateway<br>AddCache<br>AddUploadBuffer<br>AddWorkingStorage<br>CreateCachediSCSIVolume<br>CreateSnapshot<br>CreateSnapshotFromVolumeRecov-eryPoint<br>CreateStorediSCSIVolume<br>DeleteBandwidthRateLimit<br>DeleteChapCredentials<br>DeleteGateway<br>DeleteVolume<br>DescribeBandwidthRateLimit<br>DescribeCache<br>DescribeCachediSCSIVolumes<br>DescribeChapCredentials<br>DescribeGatewayInformation<br>DescribeMaintenanceStartTime<br>DescribeSnapshotSchedule<br>DescribeStorediSCSIVolumes<br>DescribeWorkingStorage<br>ListLocalDisks<br>ListGateways<br>ListVolumes<br>ListVolumeRecoveryPoints<br>ShutdownGateway<br>StartGateway<br>UpdateBandwidthRateLimit<br>UpdateChapCredentials |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| | | UpdateMaintenanceStartTime<br>UpdateGatewayInformation<br>UpdateGatewaySoftwareNow<br>UpdateSnapshotSchedule |
| `OutdatedGateway` | The specified gateway is out of date. | ActivateGateway |
| `SnapshotInProgressException` | The specified snapshot is in progress. | DeleteVolume |
| `SnapshotIdInvalid` | The specified snapshot is invalid. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume |
| `StagingAreaFull` | The staging area is full. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume |
| `TargetAlreadyExists` | The specified target already exists. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume |
| `TargetInvalid` | The specified target is invalid. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume<br>DeleteChapCredentials<br>DescribeChapCredentials<br>UpdateChapCredentials |
| `TargetNotFound` | The specified target was not found. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume<br>DeleteChapCredentials<br>DescribeChapCredentials<br>DeleteVolume<br>UpdateChapCredentials |

| Operation Error Code | Message | Operations That Return this Error Code |
|---|---|---|
| `UnsupportedOperationForGatewayType` | The specified operation is not valid for the type of the gateway. | AddCache<br>AddWorkingStorage<br>CreateCachediSCSIVolume<br>CreateSnapshotFromVolumeRecoveryPoint<br>CreateStorediSCSIVolume<br>DeleteSnapshotSchedule<br>DescribeCache<br>DescribeCachediSCSIVolumes<br>DescribeStorediSCSIVolumes<br>DescribeUploadBuffer<br>DescribeWorkingStorage<br>ListVolumeRecoveryPoints |
| `VolumeAlreadyExists` | The specified volume already exists. | CreateCachediSCSIVolume<br>CreateStorediSCSIVolume |
| `VolumeIdInvalid` | The specified volume is invalid. | DeleteVolume |
| `VolumeInUse` | The specified volume is already in use. | DeleteVolume |
| `VolumeNotFound` | The specified volume was not found. | CreateSnapshot<br>CreateSnapshotFromVolumeRecoveryPoint<br>DeleteVolume<br>DescribeCachediSCSIVolumes<br>DescribeSnapshotSchedule<br>DescribeStorediSCSIVolumes<br>UpdateSnapshotSchedule |
| `VolumeNotReady` | The specified volume is not ready. | CreateSnapshot<br>CreateSnapshotFromVolumeRecoveryPoint |

# Error Responses

When there is an error, the response header information contains:

- Content-Type: application/x-amz-json-1.1
- An appropriate `4xx` or `5xx` HTTP status code

The body of an error response contains information about the error that occurred. The following sample error response shows the output syntax of response elements common to all error responses.

```
{
    "__type": "String",
    "message": "String",
    "error":
        { "errorCode": "String",
          "errorDetails": "String"
        }
}
```

The following table explains the JSON error response fields shown in the preceding syntax.

**__type**
One of the exceptions from Exceptions (p. 383).

*Type*: String

**error**
Contains API-specific error details. In general errors (i.e., not specific to any API), this error information is not shown.

*Type*: Collection

**errorCode**
One of the operation error codes .

*Type*: String

**errorDetails**
This field is not used in the current version of the API.

*Type*: String

**message**
One of the operation error code messages.

*Type*: String

# Error Response Examples

The following JSON body is returned if you use the DescribeStorediSCSIVolumes API and specify a gateway ARN request input that does not exist.

```
{
  "__type": "InvalidGatewayRequestException",
  "message": "The specified volume was not found.",
  "error": {
    "errorCode": "VolumeNotFound"
```

```
  }
}
```

The following JSON body is returned if AWS Storage Gateway calculates a signature that does not match the signature sent with a request.

```
{
  "__type": "InvalidSignatureException",
  "message": "The request signature we calculated does not match the signature
 you provided."
}
```

# Operations in AWS Storage Gateway

For a list of AWS Storage Gateway operations, go to Actions in the *AWS Storage Gateway API Reference*.

# Appendices for AWS Storage Gateway

This appendix includes the following sections.

**Topics**


- Appendix A: Components in Your vSphere Environment for AWS Storage Gateway (p. 404)
- Appendix B: Configuring a VMware ESXi Host for AWS Storage Gateway (p. 406)
- Appendix C: Components in Your Hyper-V Environment for AWS Storage Gateway (p. 410)
- Appendix D: Configuring a Microsoft Hyper-V Host for AWS Storage Gateway (p. 411)
- Appendix E: AWS Storage Gateway Resources (p. 421)
- Appendix F: AWS Storage Gateway Limits (p. 421)
- Appendix G: Launching and Activating an Amazon EC2 AMI into a Non-Default VPC (p. 422)
- Appendix H: VTL Devices (p. 423)
- Appendix I: Testing Your Setup by Using Symantec Backup Exec 2012 (p. 426)


## Appendix A: Components in Your vSphere Environment for AWS Storage Gateway

You can use VMware to create an on-premises virtual machine to host a gateway using the AWS Storage Gateway service. You use a VMware client to interact with a VMware server and create your virtual machines. A gateway virtual machine definition—or template—that contains all the files and data for creating a new gateway is available from the AWS Storage Gateway detail page. The template is distributed as a single `.ova` file that is deployed on the VMware server. In this section, the components of the VMware vSphere environment that you need to know to use the AWS Storage Gateway service are discussed.

The following table describes the subset of vSphere components that you typically work with when using the AWS Storage Gateway service.

| Component | Description |
| --- | --- |
| VMware vSphere | The VMware virtualization platform for managing its virtual computing infrastructure including the client and server. |
| VMware ESXi hypervisor OS (vSphere Server) | The VMware server OS that hosts the gateway virtual machine. You interact with the OS through the vSphere client GUI. To provision a gateway in AWS Storage Gateway, you only need to access the host during the activation of the gateway. For all other management and maintenance-related functions, you use the AWS Management Console. |
| VMware vSphere Client (vSphere Client) | The VMware software that you use on your computer to access and manage your VMware environment. You manage your virtual machine (which contains the gateway) using the client. |
| VMware High Availability | VMware High Availability (HA) is a component of vSphere that can provide protection from failures in your infrastructure layer supporting a gateway VM. VMware HA does this by using multiple hosts configured as a cluster so that if one host running a gateway VM fails, the gateway VM can be restarted automatically on another host within the cluster. AWS Storage Gateway can be used with VMware HA. For more information about VMware HA, go to VMware HA: Concepts and Best Practices. |
| Virtual machine | The software implementation of a computer that contains the components of AWS Storage Gateway. The virtual machine (VM) runs on the VMware vSphere platform. |
| OVA, OVF | A template that represents a customized virtual machine. The AWS Storage Gateway appliance is an Open Virtualization Format (OVF) package that is distributed in an Open Virtualization Application (OVA). The OVA template contains all the information needed to configure and start a gateway. You deploy the template using the client connected to a VMware server. For instructions about downloading the OVA template for AWS Storage Gateway, go to the AWS Storage Gateway detail page. |
| Data store | The storage on the vSphere server where the files that define a virtual machine are stored. These files come from the OVA file provided as part of the service. When you deploy the OVA, you select a data store on which to store the file, if there is more than one data store for the VMware server. |

# Appendix B: Configuring a VMware ESXi Host for AWS Storage Gateway

This section provides basic information for you to set up your virtualization host. Following the basic setup, we also cover some optional host configuration.

The AWS Storage Gateway service includes an on-premises software appliance that communicates with the AWS cloud storage infrastructure. The appliance is packaged as a virtual machine that you deploy on a host running the VMware ESX/ESXi virtualization software. For more information on the VMware virtualization software, go to VMware vSphere Hypervisor. For requirements that your VMware environment must meet to run AWS Storage Gateway, see Requirements (p. 11).

### To install the VMware vSphere hypervisor OS on your host

1. Insert the VMware vSphere hypervisor disk in the disk drive.
2. Restart the computer.

   Depending on your computer bios settings, the computer might automatically boot off your disk. If not, check the relevant settings to boot the computer from the hypervisor disk.
3. Follow the instructions on the monitor to install the VMware hypervisor OS.

   This installation wipes any existing content on the disk and installs the hypervisor.

   **Tip**
   After a successful VMware hypervisor host installation, the monitor displays the IP address of the host computer. Note down this IP address. You use the IP address to connect to the host.
4. Set the time on the host.

   For instructions, see Step 2.2.3: Synchronize VM Time with Host Time (p. 20).

In the preceding steps, you provisioned a host with VMware hypervisor. The hypervisor is aware of host computer configuration, such as available processors, memory, and local hard disks. The host provides these resources to the AWS Storage Gateway.

You can optionally configure this host by adding more storage, such as additional direct-attached disks or SAN disks. The following steps illustrate how you can add one or more SAN disks to this host.

### To connect to the hypervisor host

1. Start the VMware vSphere client and connect to the host using the host IP address.

   The VMware vSphere Client dialog box appears.

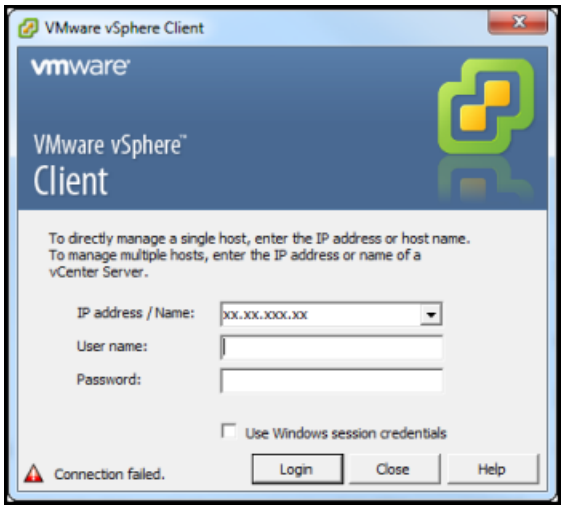


2. Enter the IP address of the host in the **IP Address** field.
3. Enter the credentials in the **User Name** and **Password** fields.
4. Click **Login**.

   This connects your client to the host. You are now ready to configure the host.

### To add a new iSCSI target

1. After you have connected to your remote device through the hypervisor, go to the **Configuration** tab of the host and click **Storage** in the **Hardware** list.

   The **Datastores** pane shows the available data stores.

   For example, the following example shows that the host has two local hard drives, datastore1, and datastore2 available.

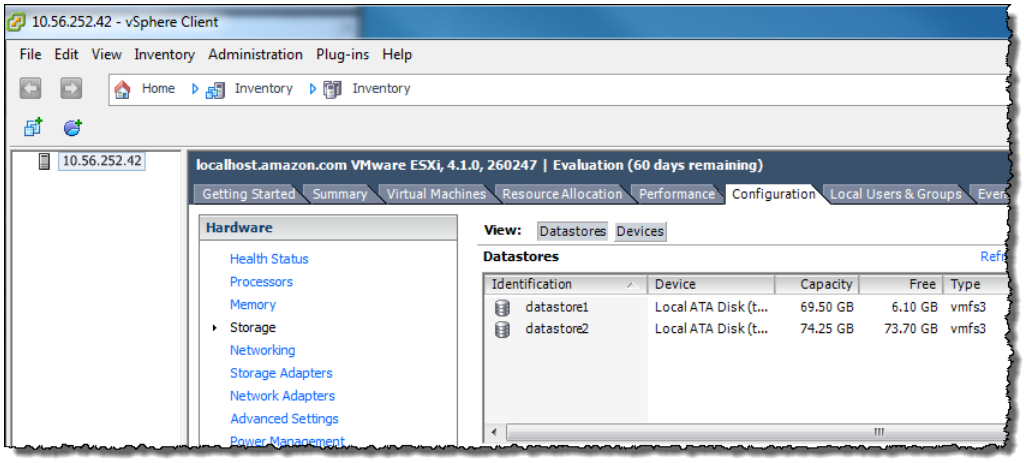


2. In the **Hardware** list, click **Storage Adapters**.
3. In the **Storage Adapters** pane, select **iSCSI Software Adapter**, and then click the **Properties** link in the **Details** pane.

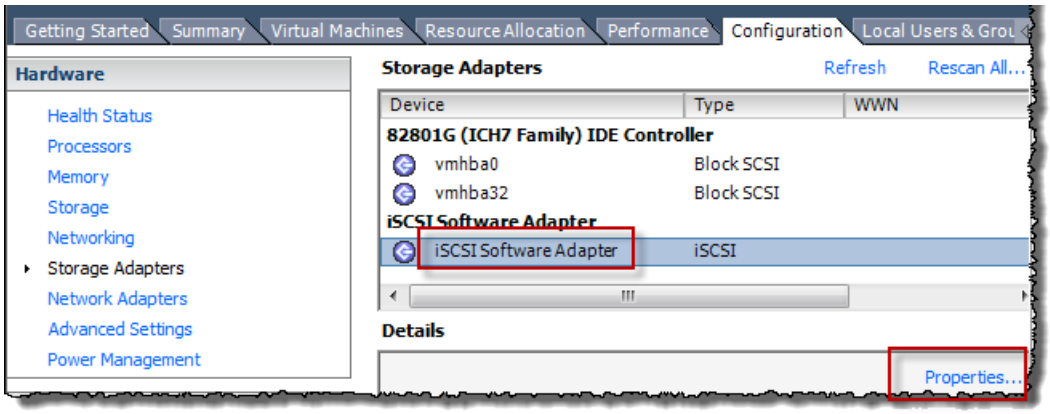


4. In the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box, click **Configure**.

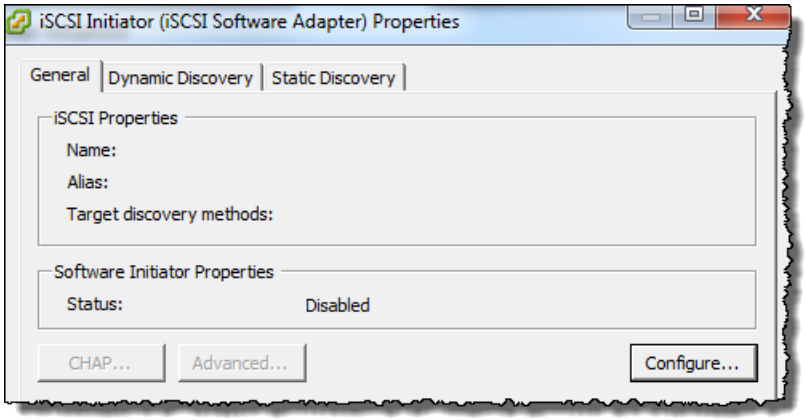


5. In the **General Properties** dialog box, select **Enabled** to set the software initiator status to enabled and click **OK**.

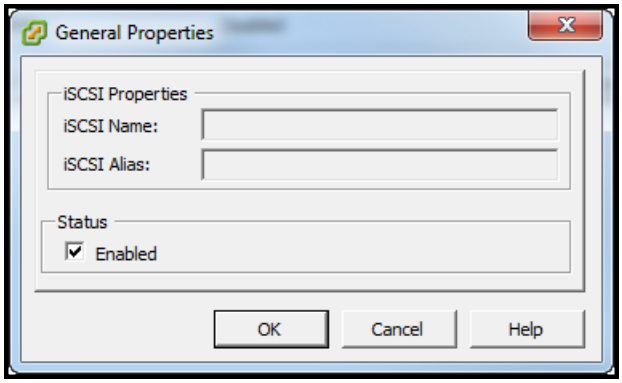


6. In the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box, select the **Dynamic Discovery** tab, and click **Add** to add an iSCSI target.

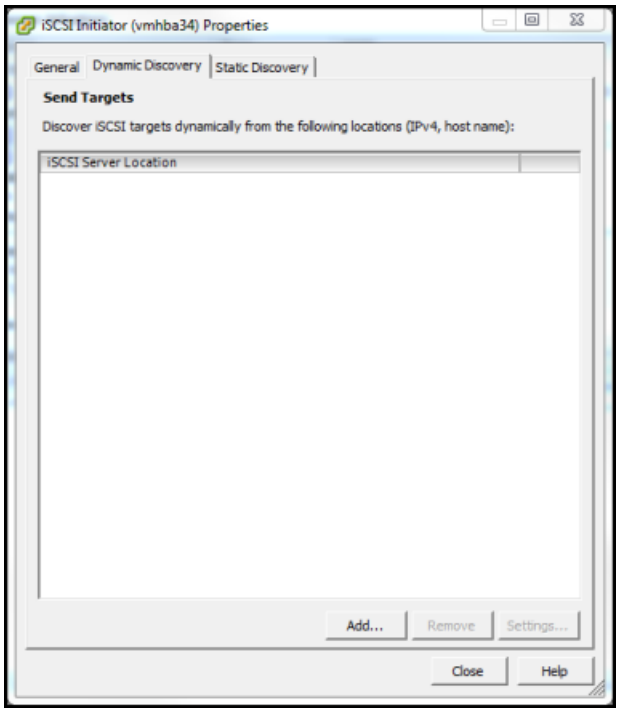


7. In the **Add Send Target Server** dialog box, enter a name in the **iSCSI Server** field and a port in the **Port** field and click **OK**.

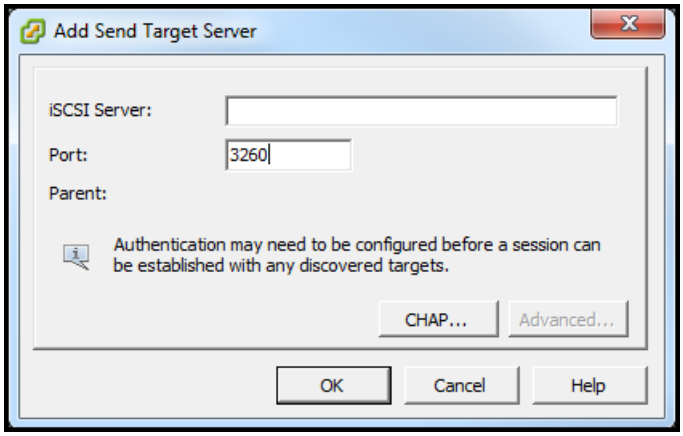


   Enter the IP address or DNS name of the storage system.

   The new iSCSI server location that is entered here appears in the **Sends Target** list on the **Dynamic Discovery** tab.

8. Click **Close** to close the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box.

At this time, you have added a new iSCSI target in the host configuration.

# Appendix C: Components in Your Hyper-V Environment for AWS Storage Gateway

You use Microsoft Hyper-V to create an on-premises virtual machine that hosts a storage gateway. You use the Hyper-V Manager to interact with a Hyper-V server and create your virtual machines. A gateway virtual machine definition—or template—that contains all the files and data for creating a new gateway is available from the AWS Storage Gateway console. The template is distributed as a single `.zip` file that you import into the Hyper-V server. This section discusses the components of the Microsoft Hyper-V environment that you need to know to use AWS Storage Gateway.

The following table describes the subset of Hyper-V components that you typically work with when using AWS Storage Gateway.

| Component | Description |
|---|---|
| Microsoft Hyper-V | The Microsoft virtualization platform for managing a virtual computing infrastructure including the client and server. |
| Hyper-V hypervisor OS | The Hyper-V server operating system (OS) that hosts the gateway virtual machine. You interact with the OS through the Microsoft Hyper-V Manager graphical user interface. To provision a storage gateway, you need to only access the host during the activation of the gateway. For all other management and maintenance-related functions, you use the AWS Management Console. |
| Hyper-V Manager | The Hyper-V client software that you use on your computer to access and manage your Hyper-V environment. You manage your virtual machine (that contains the gateway) using the client. |
| Virtual machine | The software implementation of a computer that contains the components of AWS Storage Gateway. The virtual machine (VM) runs on the Microsoft Hyper-V platform. |

| Component | Description |
| --- | --- |
| Import files (packaging of VM) | The AWS Storage Gateway appliance is distributed as a compressed directory containing the following:<br>• `Snapshots` folder, which will be empty for AWS Storage Gateway.<br>• `Virtual Hard Disks` folder, which contains one virtual hard disk file called `AWS-Storage-Gateway.vhd`.<br>**Note**<br>AWS Storage Gateway supports the .vhdx file type, which enables you create larger virtual hard disks than other file types.<br>• `Virtual Machines` folder, which contains an exported configuration files *`GUID`*`.exp`, where `GUID` is the virtual machine ID.<br>• `config.xml`, which contains configuration information used for importing.<br>**Note**<br>This `config.xml` file is only provided in the Hyper-V 2008 R2 gateway appliance.<br>You deploy AWS Storage Gateway to Hyper-V by first uncompressing the directory and then importing the uncompressed folder using the Hyper-V Manager. |

# Appendix D: Configuring a Microsoft Hyper-V Host for AWS Storage Gateway

This appendix provides basic information for you to set up, configure, and troubleshoot your Microsoft Hyper-V 2008 R2 virtualization host including:

- Setting Up and Configuring
  - Installing Microsoft Hyper-V (p. 412)
  - Connecting to Microsoft Hyper-V Host (p. 413)
  - Configuring Virtual Network Settings (p. 415)
  - Adding a Virtual Disk Backed by a Hard Disk (p. 416)
- Troubleshooting
  - Troubleshooting Your Microsoft Hyper-V Setup (p. 417)

# Setting Up and Configuring a Microsoft Hyper-V Host

The AWS Storage Gateway service includes an on-premises software appliance that communicates with the AWS cloud storage infrastructure. The appliance is packaged as a virtual machine that you can deploy on a host running Microsoft Hyper-V virtualization software. For more information on the Microsoft Hyper-V software, go to Microsoft Server Virtualization. For requirements that your Hyper-V environment must meet to run AWS Storage Gateway, see Requirements (p. 11).

## Installing Microsoft Hyper-V

The section describe a procedure for installing Microsoft Hyper-V. If you already have a Microsoft Hyper-V virtualization environment or your environment will be set up by an administrator familiar with the platform, then you do not need to understand these steps in detail.

Refer to the Hyper-V Getting Started Guide on the *Microsoft TechNet* site for more information about the installation of Hyper-V.

### To install the Microsoft Hyper-V hypervisor OS on your host

1. Insert the Microsoft Hyper-V disk in the disk drive.
2. Restart the computer.

   Depending on your computer's BIOS settings, the computer might automatically boot off your disk. If not, check the relevant settings to boot the computer from the hypervisor disk.
3. Follow the instructions on the monitor to install the Hyper-V hypervisor OS.

   This installation wipes any existing content on the disk and installs the hypervisor.

   After a successful Hyper-V hypervisor host installation, you will be prompted to create an Administrator account password. After creating this account, the monitor displays a **Server Configuration** menu where you will do further configuration of the host.
4. In the **Server Configuration** menu, configure the host. We recommend the following:

| To... | Do This... |
|---|---|
| Configure remote management. | Select option 4 and then enable the following:<br>• option 1, **Allow MMC Remote Management**<br>• option 2, **Enable Windows PowerShell**<br>• option 3, **Allow Server Manager Remote Management** |
| Find the network address of the host. | Select option 8 and follow the prompts. Note the IP address for use later. |
| Set the date and time. | Select option 9 and follow the prompts. |
| (Optional) Change the computer name. | Select option 2 and follow the prompts. Since this requires a reboot, you may want to make this configuration change last. |
| (Optional) Add local administrators. | Select option 3 and follow the prompts. |

| To... | Do This... |
|---|---|
| (Optional) Enable remote desktop. | Select option 7 and follow the prompts. |

The following example shows a Server Configuration menu.

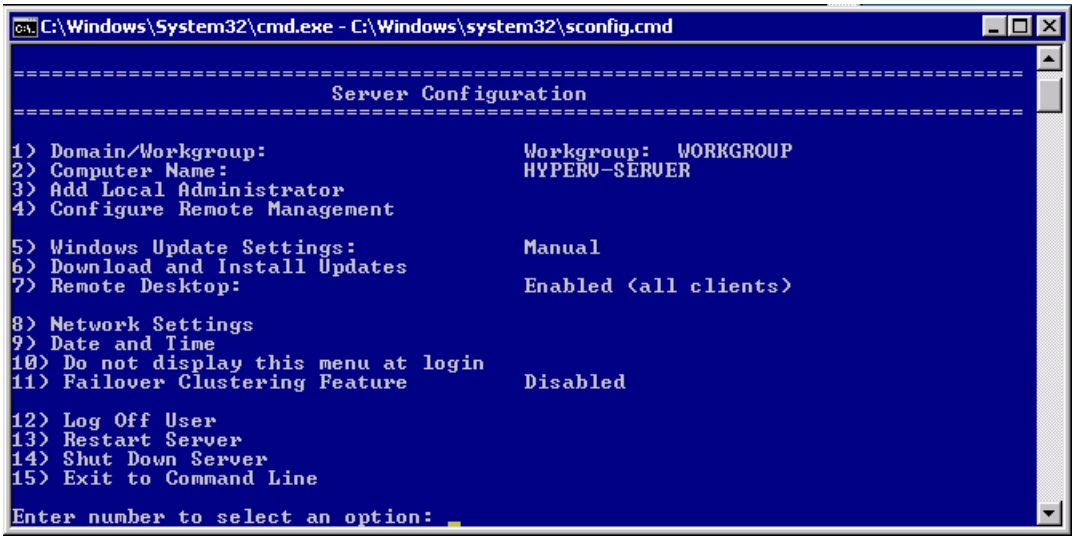
C:\Windows\System32\cmd.exe - C:\Windows\system32\sconfig.cmd
Server Configuration
1) Domain/Workgroup: Workgroup: WORKGROUP
2) Computer Name: HYPERV-SERVER
3) Add Local Administrator
4) Configure Remote Management
5) Windows Update Settings: Manual
6) Download and Install Updates
7) Remote Desktop: Enabled (all clients)
8) Network Settings
9) Date and Time
10) Do not display this menu at login
11) Failover Clustering Feature Disabled
12) Log Off User
13) Restart Server
14) Shut Down Server
15) Exit to Command Line
Enter number to select an option:

5. (Optional) You may need to put the IP address of the hypervisor host in your hosts file of client computers that connect to the hypervisor host.

   For example, in Windows 7 and 8, the hosts file can be found at this location:

```
%SystemRoot%\system32\drivers\etc\hosts
```

## Connecting to Microsoft Hyper-V Host

The Hyper-V Manager runs on your client computer and connects to the hypervisor host. You use the Microsoft Hyper-V Manager to import, configure, and start the AWS Storage Gateway VM.

### To connect to the hypervisor host

1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).

   **Note**
   The Hyper-V Manager is a feature that you enable for your client computer. For more information about enabling it, go to Install and Configure Hyper-V Tools for Remote Administration.

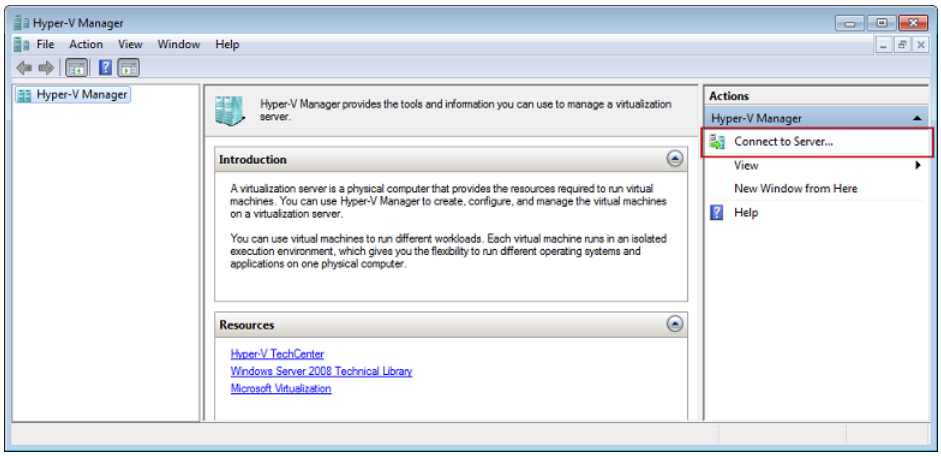
Hyper-V Manager
File Action View Window Help
Hyper-V Manager
Hyper-V Manager provides the tools and information you can use to manage a virtualization server.
Introduction
A virtualization server is a physical computer that provides the resources required to run virtual machines. You can use Hyper-V Manager to create, configure, and manage the virtual machines on a virtualization server.
You can use virtual machines to run different workloads. Each virtual machine runs in an isolated execution environment, which gives you the flexibility to run different operating systems and applications on one physical computer.
Resources
Hyper-V TechCenter
Windows Server 2008 Technical Library
Microsoft Virtualization
Actions
Hyper-V Manager
Connect to Server...
View
New Window from Here
Help

2. In the **Actions** pane, select **Connect to Server...**.
3. In the **Select Computer** dialog box, select **Another computer** and either type the IP address of the hypervisor host or the hostname and click **OK**.

   **Note**
   To connect to a hypervisor host using the hostname, you may need to make an entry in your `hosts` file so that the hostname can be mapped to the correct IP address.

   **Note**
   If you have not been added to the local administrators for the hypervisor host, you may be prompted for credentials.

The following example shows Hyper-V Manager connected to a hypervisor host called `HYPERV-SERVER` with one gateway VM.

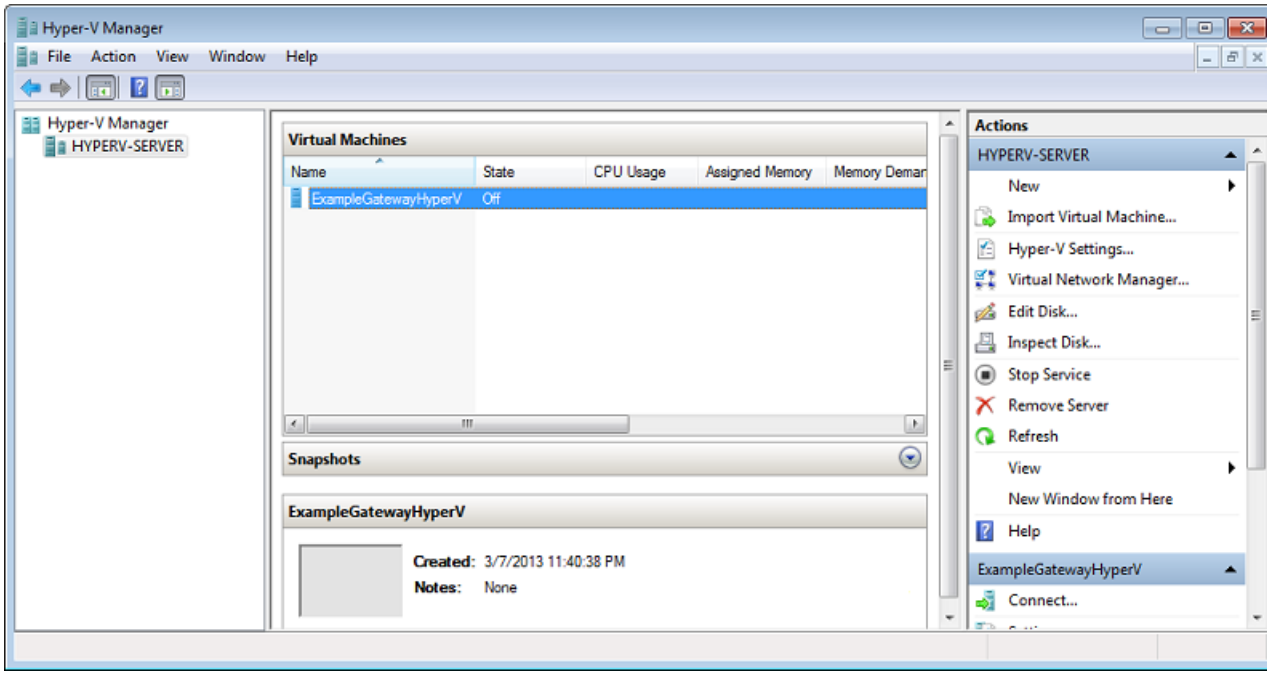

# Configuring Virtual Network Settings

After you install and configure a Microsoft Hyper-V host, we recommend that you set up virtual networks by creating a new virtual network and associating it with a network interface of the host. Later, when you configure your gateway VM, you must associate it with one or more virtual networks so that the VM has connectivity.

### To configure virtual network settings for the your Hyper-V host

1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).
2. In the hypervisor host list (left pane), select your hypervisor.
3. In the **Actions** menu, under the hypervisor host name (e.g., `HYPERV-SERVER`), click **Virtual Network Manager**.

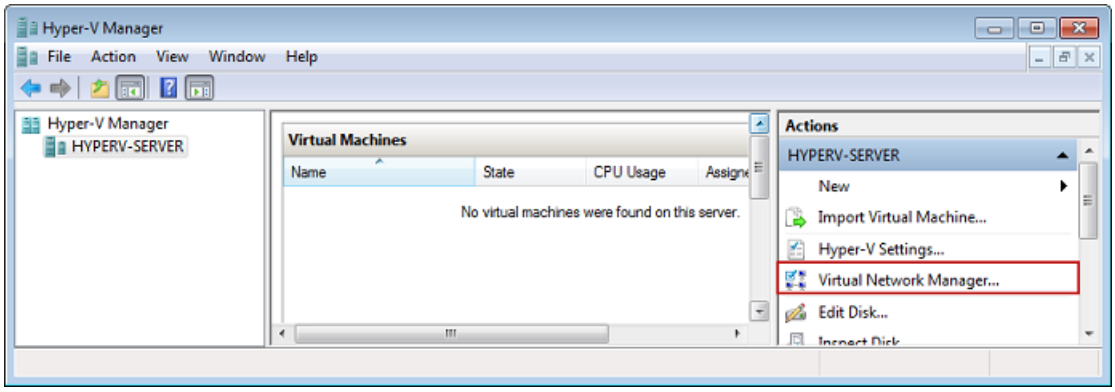


4. In the **Virtual Network Manager** dialog box, select **New virtual network**.

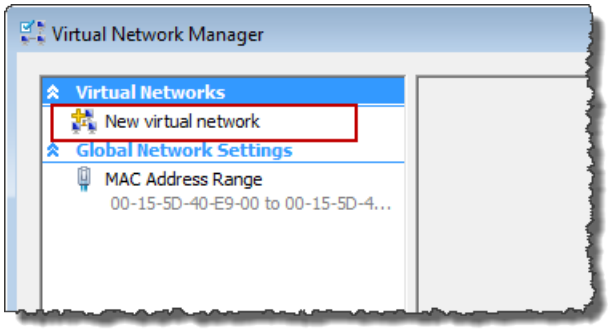


5. Select **External** as the virtual network type and click **Add**.

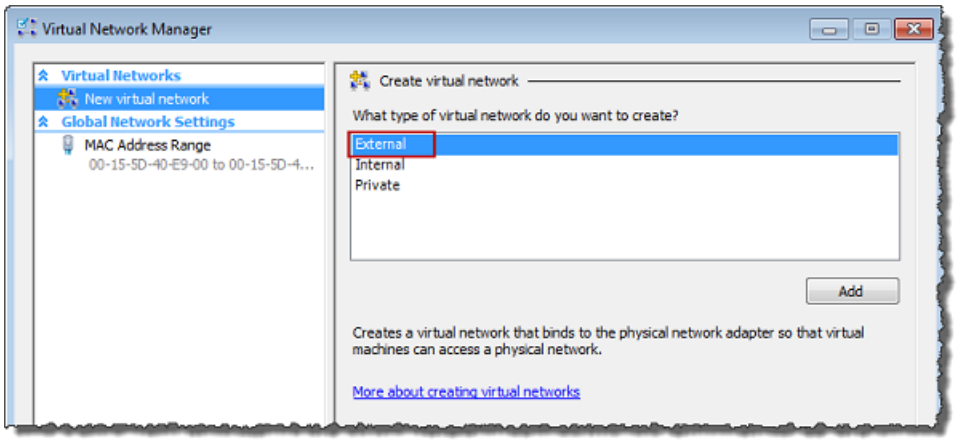
Virtual Network Manager
Virtual Networks
New virtual network
Global Network Settings
MAC Address Range
00-15-5D-40-E9-00 to 00-15-5D-4...
Create virtual network
What type of virtual network do you want to create?
External
Internal
Private
Add
Creates a virtual network that binds to the physical network adapter so that virtual machines can access a physical network.
More about creating virtual networks

6. Provide a name for the network, and click **OK**.

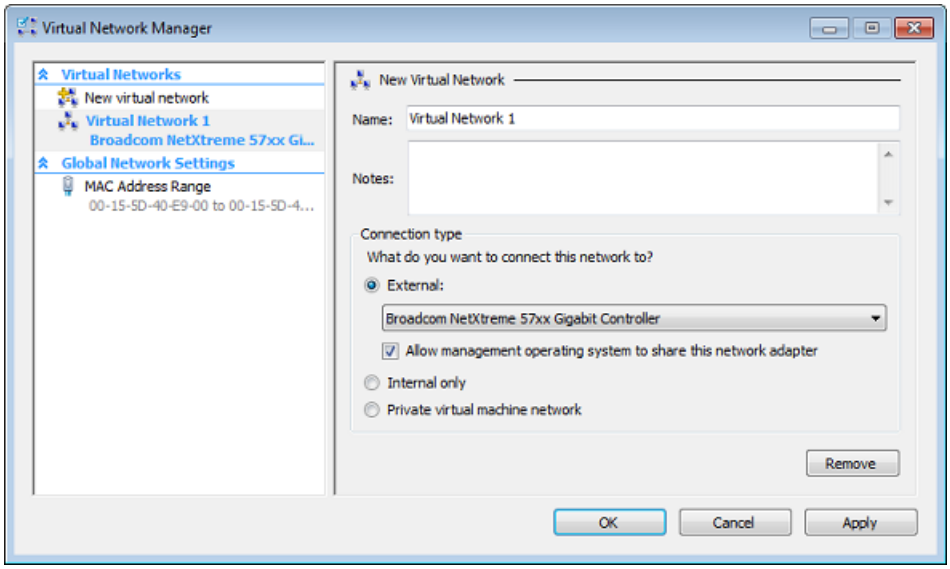


When you configure your gateway virtual machine, you can use this virtual network.

## Adding a Virtual Disk Backed by a Hard Disk

In a preceding section, you provisioned a host with Hyper-V hypervisor. The hypervisor is aware of host computer configuration, such as available processors, memory, and local hard disks. The host provides these resources to AWS Storage Gateway. You can optionally configure this host by adding more storage, such as additional direct-attached disks or SAN disks. In this section, we show you how to add a virtual disk backed by a direct-attached disk.

### To add a virtual disk backed by a physical hard disk

1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).
2. Select the VM.
3. In the **Actions** list for the VM, click **Settings...**.
4. In the **Hardware** list, click **SCSI Controller**.
5. Select **Hard Drive** in the **SCSI Controller** pane and click **Add**.

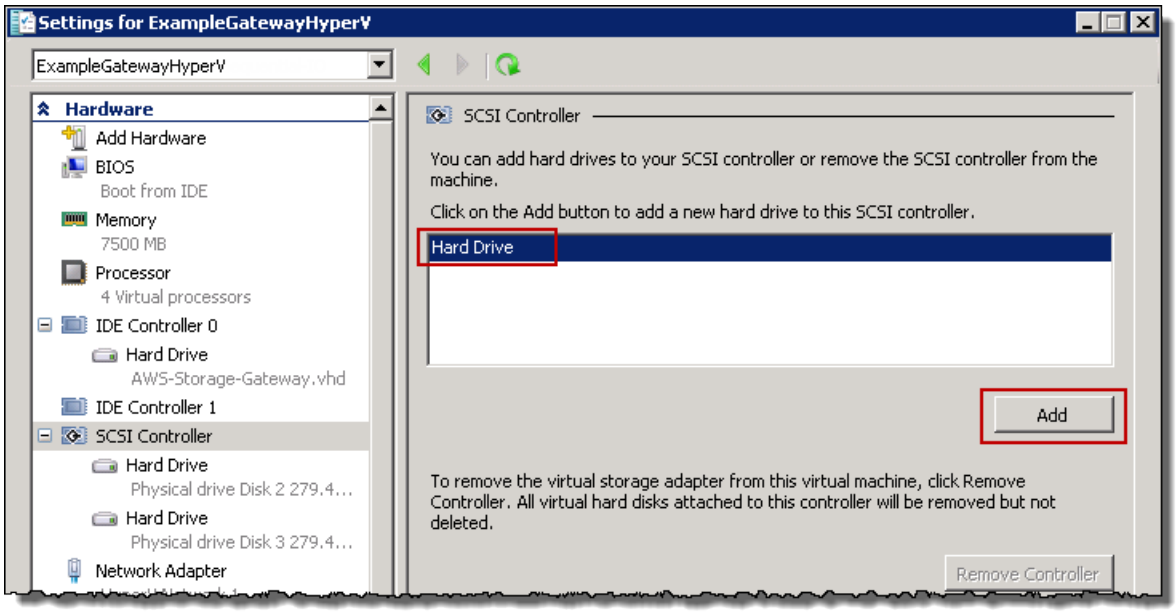
Settings for ExampleGatewayHyperV
ExampleGatewayHyperV
Hardware
Add Hardware
BIOS
Boot from IDE
Memory
7500 MB
Processor
4 Virtual processors
IDE Controller 0
Hard Drive
AWS-Storage-Gateway.vhd
IDE Controller 1
SCSI Controller
Hard Drive
Physical drive Disk 2 279.4...
Hard Drive
Physical drive Disk 3 279.4...
Network Adapter
SCSI Controller
You can add hard drives to your SCSI controller or remove the SCSI controller from the machine.
Click on the Add button to add a new hard drive to this SCSI controller.
Hard Drive
Add
To remove the virtual storage adapter from this virtual machine, click Remove Controller. All virtual hard disks attached to this controller will be removed but not deleted.
Remove Controller

6. In the **Hard Drive** pane, select **Physical hard disk**.

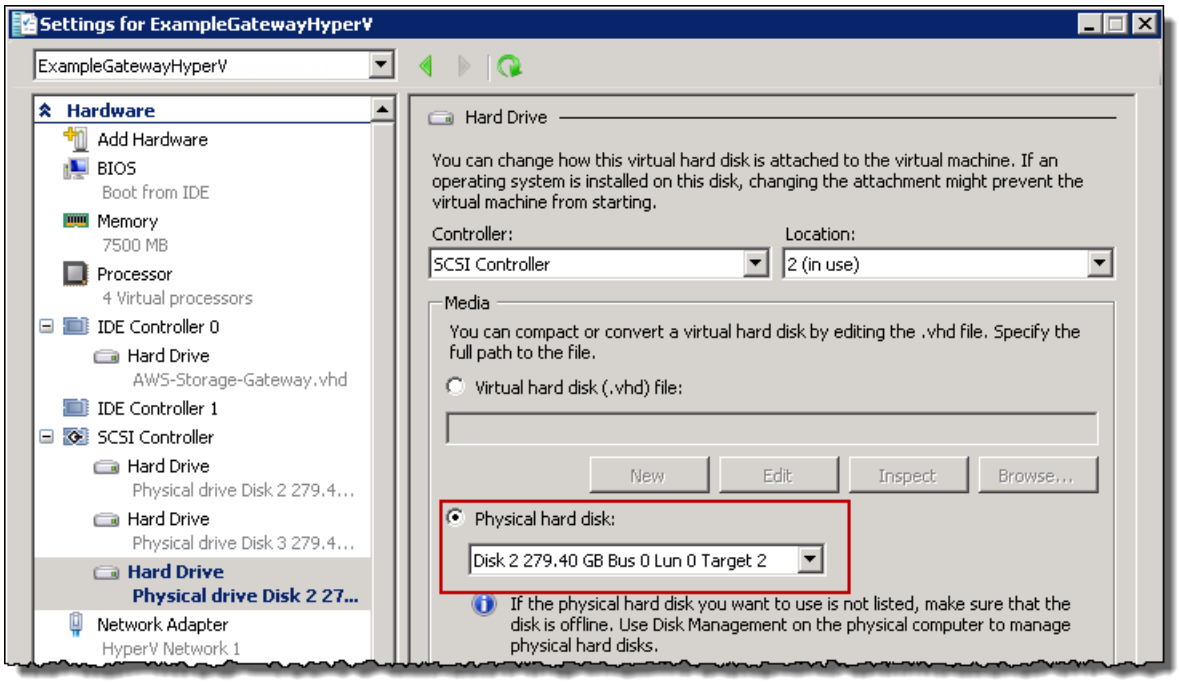
Settings for ExampleGatewayHyperV
ExampleGatewayHyperV
Hardware
Add Hardware
BIOS
Boot from IDE
Memory
7500 MB
Processor
4 Virtual processors
IDE Controller 0
Hard Drive
AWS-Storage-Gateway.vhd
IDE Controller 1
SCSI Controller
Hard Drive
Physical drive Disk 2 279.4...
Hard Drive
Physical drive Disk 3 279.4...
Hard Drive
Physical drive Disk 2 27...
Network Adapter
HyperV Network 1
Hard Drive
You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.
Controller:
SCSI Controller
Location:
2 (in use)
Media
You can compact or convert a virtual hard disk by editing the .vhd file. Specify the full path to the file.
Virtual hard disk (.vhd) file:
New
Edit
Inspect
Browse...
Physical hard disk:
Disk 2 279.40 GB Bus 0 Lun 0 Target 2
If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.

7. Click **OK**.

# Troubleshooting Your Microsoft Hyper-V Setup

The following table lists typical issues that you might encounter when deploying AWS Storage Gateway on the Microsoft Hyper-V platform.

| Issue | Action to Take |
| --- | --- |
| You try to import a gateway and receive the error message: "Import failed. Unable to find virtual machine import file under location ...".<br> 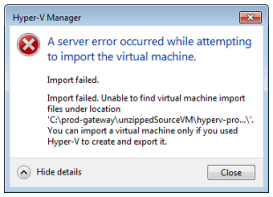 | This error can occur for the following reasons:<br>• If you are not pointing to the root of the unzipped gateway source files. The last part of the location you specify in the **Import Virtual Machine** dialog box should be `AWS-Storage-Gateway`, as the following example shows:<br> 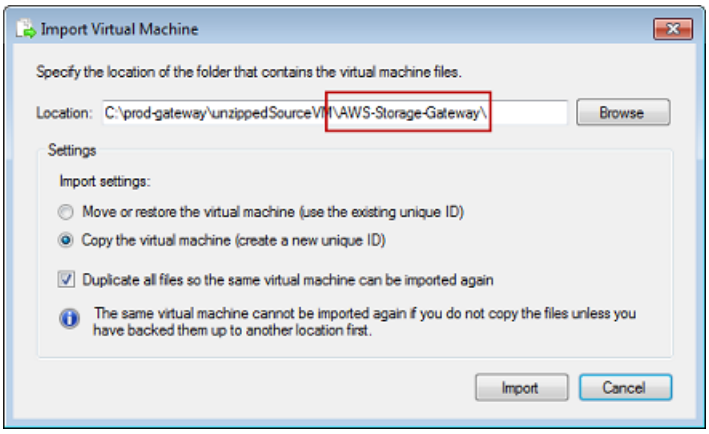 <br>• If you have already deployed a gateway and you did not select the **Copy the virtual machine** option and check the **Duplicate all files...** option in the **Import Virtual Machine** dialog box, then the VM was created in the location where you have the unzipped gateway files and you can not import from this location again. To fix this problem, get a fresh copy of the unzipped gateway source files and copy to a new location. Use the new location as the source of the import. The following example shows the options that you must check if you plan on creating multiple gateways from one unzipped source files location.<br> 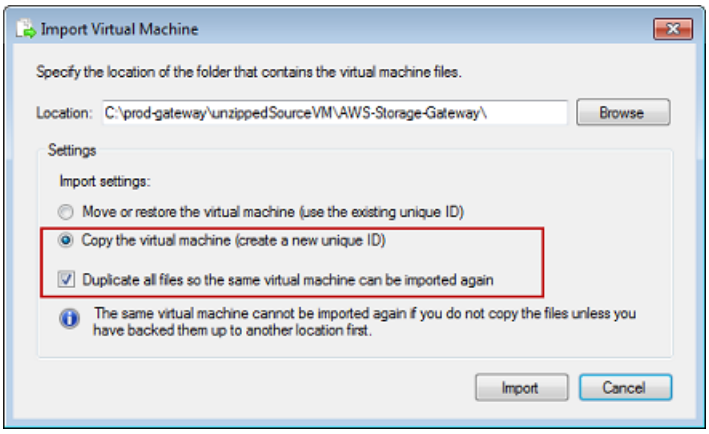 <br>For more information about deploying the gateway, see Step 2.2: Download and Deploy the AWS Storage Gateway VM on Your Host (p. 36) in the Getting Started exercise for Microsoft Hyper-V. |

| Issue | Action to Take |
| --- | --- |
| You try to import a gateway and receive the error message: "Import failed. Import task failed to copy file."<br> 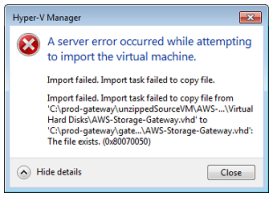 | If you have already deployed a gateway and you try to reuse the default folders that store the virtual hard disk files and virtual machine configuration files, then this error will occur. To fix this problem, specify new locations in the **Hyper-V Settings** dialog box.<br> 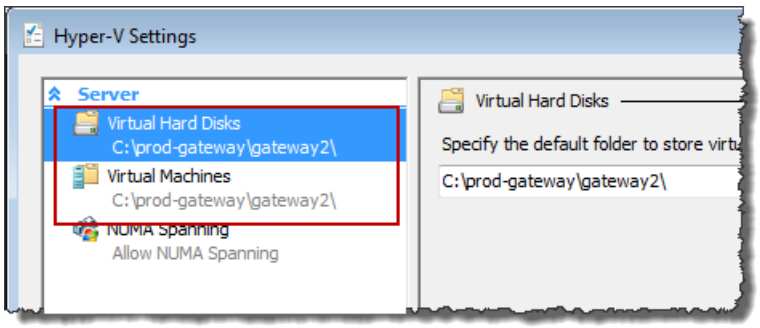 <br>For more information about deploying the gateway, see To import the VM (p. 39) in the exercise for Microsoft Hyper-V. |
| You try to import a gateway and receive an error message: "Import failed. Import failed because the virtual machine must have a new identifier. Select a new identifier and try the import again."<br> 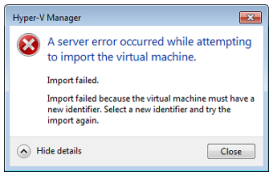 | When you import the gateway make sure you select the **Copy the virtual machine** option and check the **Duplicate all files...** option in the **Import Virtual Machine** dialog box to create a new unique ID for the VM. The following example shows the options in the **Import Virtual Machine** dialog box that you should use.<br> 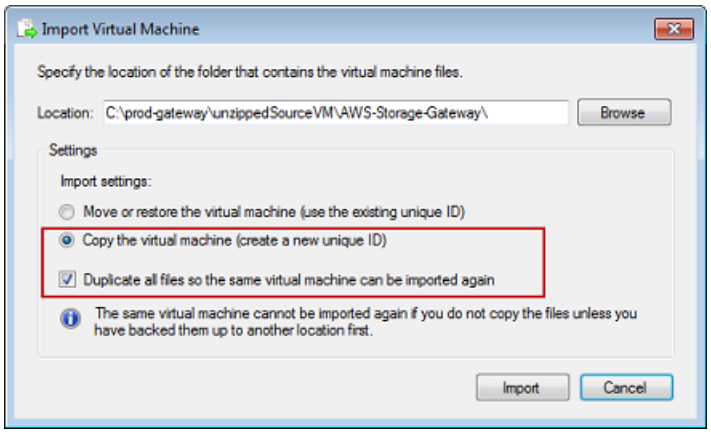 <br>For more information about importing the gateway, see To import the VM (p. 39) in the exercise for Microsoft Hyper-V. |

| Issue | Action to Take |
| --- | --- |
| You try to start a gateway VM and receive an error message "The child partition processor setting is incompatible with parent partition."<br> 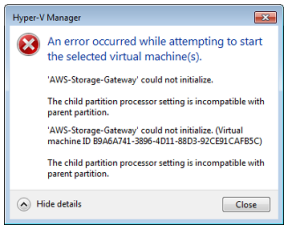 | This error is likely caused by a CPU discrepancy between the required CPUs for the gateway and the available CPUs on the host. Ensure that the VM CPU count is supported by the underlying hypervisor.<br>For more information about the requirements for AWS Storage Gateway, see Requirements (p. 11). |
| You try to start a gateway VM and receive an error message "Failed to create partition: Insufficient resources exist to complete the requested service."<br> 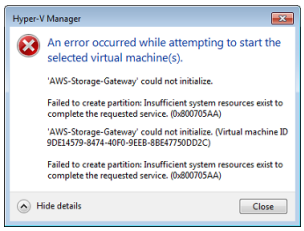 | This error is likely caused by a RAM discrepancy between the required RAM for the gateway and the available RAM on the host.<br><br>For more information about the requirements for AWS Storage Gateway, see Requirements (p. 11). |
| Your snapshots and gateway software updates are occurring at slightly different times than expected. | The gateway VM's clock may be offset from the actual time, known as clock drift. Check and correct the VM's time using local gateway console's time synchronization option. For more information, see Synchronizing Your Gateway VM Time (p. 146). |
| You need to put the unzipped Microsoft Hyper-V AWS Storage gateway files on the host file system. | Access the host as you would a typical Microsoft Windows server. For example, if the hypervisor host is name `hyperv-server`, then you can use the following UNC path `\\hyperv-server\c$`, which assumes that the name `hyperv-server` can be resolved or is defined in your local hosts file. |
| You are prompted for credentials when connecting to hypervisor.<br> 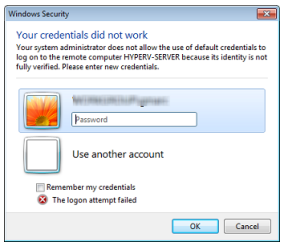 | Add your user credentials as a local administrator for the hypervisor host by using the Sconfig.cmd tool. For more information, see Setting Up and Configuring a Microsoft Hyper-V Host (p. 412). |

# Appendix E: AWS Storage Gateway Resources

The source code for certain open source software components that are included with the AWS Storage Gateway software is available for download at:

- https://s3.amazonaws.com/aws-storage-gateway-terms/sources.tar for gateways deployed on VMware ESXi
- https://s3.amazonaws.com/aws-storage-gateway-terms/sources_hyperv.tar for gateways deployed on Microsoft Hyper-V

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/).

The packages comprising the AWS Storage Gateway VM are tracked and monitored for security vulnerabilities. When updates are issued, they are applied to each gateway and the updated packages will increment their version number, although the major version number of the Linux distribution may not increment. For more information about managing updates, see Managing Gateway Updates Using the AWS Storage Gateway Console (p. 129).

# Appendix F: AWS Storage Gateway Limits

The following table lists limits for volumes, virtual tapes, and upload and download rates in AWS Storage Gateway.

| | Gateway-Cached Volumes | Gateway-Stored Volumes | Gateway–Virtual Tape Libraries |
|---|---|---|---|
| Maximum size of a volume | 32 TiB | 16 TiB | — |
| Total maximum size of all volumes for a gateway | 150 TiB | 192 TiB | — |
| Maximum size of a cache storage disk | 16 TiB | — | 16 TiB |
| Total maximum size of all cache disks for a gateway | 16 TiB | — | 16 TiB |
| Maximum size of an upload buffer disk | 2 TiB | 2 TiB | 2 TiB |
| Total maximum size of all upload buffer disks for a gateway | 2 TiB | 2 TiB | 2 TiB |
| Minimum download rate | 51,200 bps | 51,200 bps | — |
| Minimum upload rate | 102,400 bps | 102,400 bps | — |
| Maximum size of a virtual tape | — | — | 2.5 TiB |

| | Gateway-Cached Volumes | Gateway-Stored Volumes | Gateway–Virtual Tape Libraries |
|---|---|---|---|
| Minimum size of a virtual tape | — | — | 100 GiB |
| Total size of all tapes in a gateway | — | — | 150 TiB |
| Maximum number of volumes for a gateway | 20 | 12 | — |
| Maximum number of virtual tapes for a gateway | — | — | 1500 |

# Appendix G: Launching and Activating an Amazon EC2 AMI into a Non-Default VPC

You can launch and activate your gateway ami into a default VPC (Virtual Private Network) or nondefault VPC. This section lists the steps you follow to create a nondefault VPC for your gateway ami. To learn about VPC, go to What is Amazon VPC?.

### To launch and activate an Amazon EC2 gateway ami into a nondefault VPC

1. Set Up the VPC and Internet Gateway. Take note of your VPC id. For information about how to create a VPC, go to Getting Started with Amazon VPC.

   A subnet and an Igw (Internet gateway) are automatically created for your VPC. The subnet is associated with the Igw.
2. A routing table is created for your VPC and is associated with the Igw.
3. Create security group rules for your VPC. For more information, go to Step 3: Set Up a Security Group for Your VPC in the VPC Getting Started Guide.
4. Create a network access control list (ACL) for your VPC.

   For inbound traffic create a rule to allow port 80 (HTTP) and set the source to 0.0.0.0/0.

   For inbound traffic create a rule to allow ephemeral ports. For more information, go to Ephemeral Ports in the VPC User Guide.

   For outbound traffic create a rule to allow port 443 (HTTPS), port 1024- 65535 and set the destination 0.0.0.0/0 and then associate the ACL with your subnet.

   For more information, go to Network ACLs in the VPC User Guide.
5. Launch your gateway ami into the VPC.

   **Note**
   You must assign a public IP address to your instance. For more information, go to Assigning a Public IP Address During Launch in the VPC User Guide.

| To... | Click here. |
|---|---|
| **Launch your gateway-cached ami into a VPC.** | Amazon EC2 Gateway (p. 57) |

| To... | Click here. |
|---|---|
| **Launch your gateway-VTL ami into a VPC.** | Step 2: Deploy and Activate a Gateway on Amazon EC2 (p. 58) |

# Appendix H: VTL Devices

Gateway-VTL provides the following SCSI devices, which you select when activating your gateway.

### Medium Changer

- AWS-Gateway-VTL
- STK-L700—This device is emulated.

**Important**
The type of medium changer you select depends on the backup software you plan to use. The following table shows which medium changer to select for your backup software.

| Backup Software | Medium Changer Type |
|---|---|
| Symantec NetBackup Version 7.x | `AWS-Gateway-VTL` |
| Symantec Backup Exec 2012 | `STK-L700` |
| Symantec Backup Exec 2014 | `AWS-Gateway-VTL` |
| Symantec Backup Exec 15 | `AWS-Gateway-VTL` |
| Microsoft System Center 2012 R2 Data Protection Manager | `STK-L700` |
| Veeam Backup & Replication V7 | `STK-L700` |
| Veeam Backup & Replication V8 | `STK-L700` |
| Dell NetVault Backup 10.0 | `STK-L700` |

**Note**
You must select the medium changer that is recommended for your backup software. Other medium changers might not function properly. You can select a different medium changer after the gateway is activated. For more information, see Selecting a Medium Changer After Gateway Activation (p. 424).

### Tape Drive

- IBM-ULT3680-TD5—This device is emulated.

The following screenshot shows the activation page for a gateway-VTL.

AWS Storage Gateway
Activating Your AWS Storage Gateway Virtual Machine
Below is the type of the gateway you are activating:
Gateway Configuration: Gateway-Virtual Tape Library
Activated gateways are billed at $125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply. The AWS Service Terms are available here.
Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.
AWS Region:
US West (Oregon)
Gateway Time Zone:
(GMT) Western Europe Time, London, Lisbon, Casablanca
Gateway Name:
MyGateway-VTL
Medium Changer Type:
AWS-Gateway-VTL
Learn More
Tape Drive Type:
IBM-ULT3580-TD5
Activate My Storage Gateway
Click here if you need to exit the activation process.

# Selecting a Medium Changer After Gateway Activation

When you activate your gateway-VTL, you select a medium changer type for the gateway. You can select a different medium changer type after the gateway is activated.

> **Important**
> If your gateway-VTL uses the Symantec Backup Exec 2014 or Netbackup 7.x backup software, you must select the AWS-Gateway-VTL device type. For more information on how to change the medium changer after gateway activation for these applications, go to Best Practices for using Symantec Backup products (NetBackup, Backup Exec) with the Amazon Web Services (AWS) Storage Gateway-VTL in *Symantec Support*.

### To select a different medium changer type after gateway activation

1. Stop any related jobs that are running in your backup software.
2. On the Windows server, open the iSCSI initiator properties window.
3. Click the **Targets** tab to display the discovered targets.
4. On the Discovered targets pane, click the medium changer you want to change, click **Disconnect**, and then click **OK**.
5. On the Navigation pane of the AWS Storage Gateway console, click the gateway-VTL for which you want to change the medium changer.
6. Click the **VTL Devices** tab, and then click the medium changer in the **Tape Device ID** column.
7. Click the **Details** tab, and then click the **Configure Device Type** link in the **VTL Device Type** row.

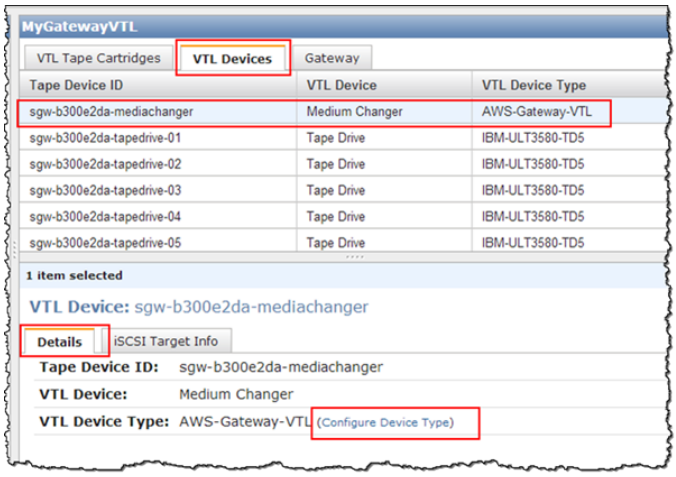


8. On the **Configure Device Type** page, click the **Device Type** box, select a medium changer, and then click **Save**.

# Updating the Device Driver for Your Medium Changer

Depending on the backup software you use on your Windows server, you might need to update the driver for your medium changer.

1. Open Device Manager on your Windows server, and expand the **Medium Changer devices** tree.
2. Right-click **Unknown Medium Changer**, and click **Update Driver Software** to open the **Update Driver Software-unknown Medium Changer** window.
3. 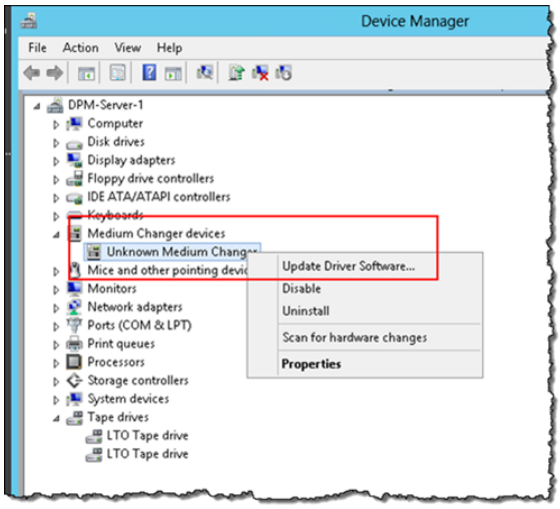


   In the **How do you want to search for driver software?** section, click **Browse my computer for driver software**.
4. Click **Let me pick from a list of device drivers on my computer**.

   **Note**
   We recommend using the Sony - TSL-A500C Autoloader driver with the Veeam Backup & Replication V7 and Microsoft System Center 2012 R2 Data Protection Manager backup software. This Sony driver has been tested with both backup software.

5. In the **Select the device driver you want to install for this hardware** section, click to uncheck the **Show compatible hardware** check box, click **Sony** in the **Manufacturer** list, click **Sony - TSL-A500C Autoloader** in the **Model** list, and then click **Next**.

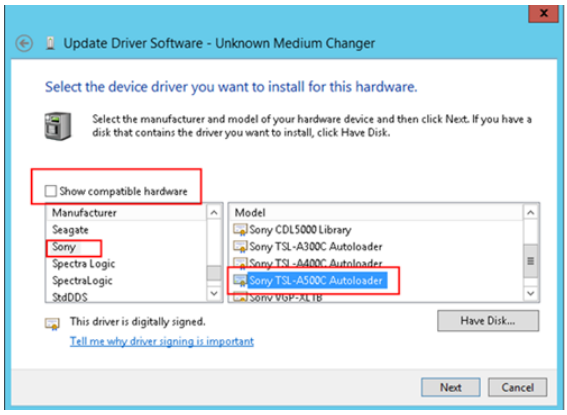
Update Driver Software - Unknown Medium Changer
Select the device driver you want to install for this hardware.
Select the manufacturer and model of your hardware device and then click Next. If you have a disk that contains the driver you want to install, click Have Disk.
Show compatible hardware
Manufacturer
Seagate
Sony
Spectra Logic
SpectraLogic
StdDDS
Model
Sony CDL5000 Library
Sony TSL-A300C Autoloader
Sony TSL-A400C Autoloader
Sony TSL-A500C Autoloader
This driver is digitally signed.
Tell me why driver signing is important
Have Disk...
Next
Cancel

6. In the warning box that appears, click **Yes**. If the driver is successfully installed, close the **Update drive software** window.

# Appendix I: Testing Your Setup by Using Symantec Backup Exec 2012

AWS Storage Gateway-VTL supports backing up data to virtual tapes and archiving the tapes in the virtual tape shelf (VTS) by using Symantec Backup Exec 2012. In this topic, we provide basic documentation needed to perform this type of backups. For detailed information about how to use Backup Exec 2012, go to the Backup Exec documentation on the Symantec website. This topic shows you how to configure storage, write data to a virtual tape, archive the tape in the VTS, and restore the data. AWS Storage Gateway-VTL supports newer versions of Symantec Backup Exec. For more information, see Testing Your Setup by Using Symantec Backup Exec (p. 263).

**Topics**


- Configuring Storage in Backup Exec 2012 (p. 426)
- Importing a Tape in Backup Exec 2012 (p. 427)
- Writing Data to a Tape in Backup Exec 2012 (p. 428)
- Archiving a Tape Using Backup Exec 2012 (p. 428)
- Restoring Data from a Tape Archived in Backup Exec 2012 (p. 429)
- Disabling a Tape Drive in Backup Exec 2012 (p. 429)


## Configuring Storage in Backup Exec 2012

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Backup Exec 2012 storage to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 4: Connect Your Gateway-VTL Devices to Your Windows Client (p. 245).

**To configure storage**

1. Start the Backup Exec 2012 software, and then click the yellow icon in the top left corner on the toolbar.
2. Click **Configuration and Settings**, and then click **Backup Exec Services** to open the Backup Exec Service Manager.

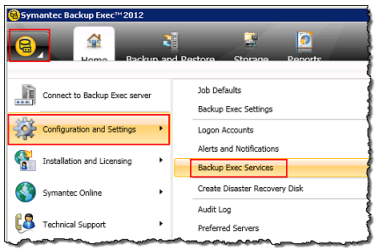


3. Click **Restart All Services**. Backup Exec 2012 recognizes the VTL devices (that is, the medium changer and tape drives). The restart process might take a few minutes.

   **Note**
   AWS Storage Gateway-VTL provides 10 tape drives. However, your Backup Exec 2012 license agreement might require your backup software to work with fewer than 10 tape drives. In that case, you must disable tape drives in the Backup Exec 2012 robotic library to leave only the number of tape drives allowed by your license agreement enabled. For instructions, see Disabling a Tape Drive in Backup Exec 2012 (p. 429).

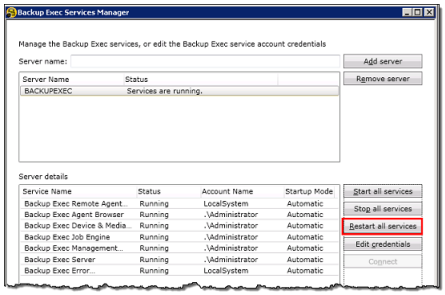


4. After the restart is completed, close the Backup Exec Service Manager.

# Importing a Tape in Backup Exec 2012

You are now ready to import a tape from your gateway into a slot.

1. Click the **Storage** tab, and then expand the **Robotic library** tree to display the VTL devices.

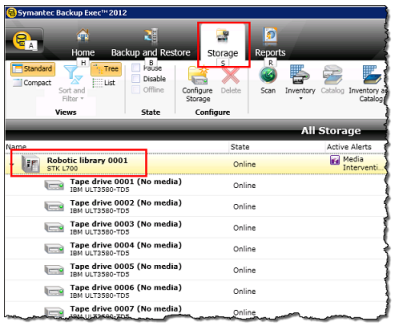


2. Double-click the **Slots** icon to display all slots.

   **Note**
   When you import tapes into the robotic library, the tapes are stored in slots instead of tape drives. Therefore, the tape drives might have a message that indicates there is no media in the drives (No media). When you initiate a backup or restore job, the tapes will be moved into the tape drives.
   You must have tapes available in your gateway tape library to import a tape into a storage slot. For instructions on how to create tapes, see Creating Virtual Tapes in Your Gateway-VTL (p. 279).

3. Right-click an empty slot, select **Import media now**, and then click **Import media now** in the submenu. In the following screenshot, slot number **3** is empty. You can select more than one slot and import multiple tapes in a single import operation.

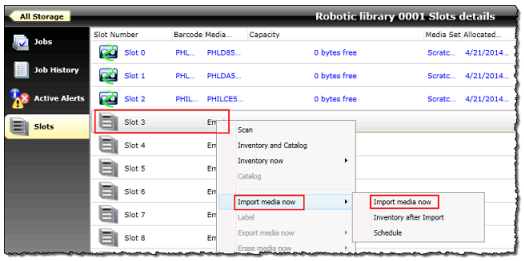


4. In the **Media Request** window that appears, click on the **Details** link.

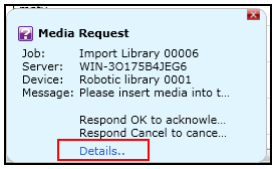


5. In the **Action Alert: Media Intervention** window, click **Respond OK** to insert the media into the slot.

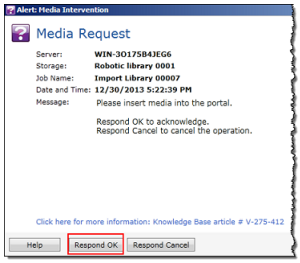
Alert: Media Intervention
Media Request
Server: WIN-3O175B4JEG6
Storage: Robotic library 0001
Job Name: Import Library 00007
Date and Time: 12/30/2013 5:22:39 PM
Message: Please insert media into the portal.
Respond OK to acknowledge.
Respond Cancel to cancel the operation.
Click here for more information: Knowledge Base article # V-275-412
Help
Respond OK
Respond Cancel

The tape appears in the slot you selected.

> **Note**
> Tapes that are imported include empty tapes and tapes that have been retrieved from the VTS to the gateway.

# Writing Data to a Tape in Backup Exec 2012

You write data to a gateway-VTL virtual tape by using the same procedure and backup policies you do with physical tapes. For detailed information, go to the Backup Exec documentation on the Symantec website.

# Archiving a Tape Using Backup Exec 2012

When you archive a tape, AWS Storage Gateway moves the tape from your gateway's virtual tape library (VTL) to the VTS—that is, the offline storage. You begin tape archival by exporting the tape using your backup software—that is, Backup Exec 2012.

### To archive your tape in VTS

1. Click the **Storage** menu, double-click **Slots**, right-click the slot you want to export the tape from, click **Export media now**, and then click **Export media now**. You can select more than one slot and export multiple tapes in a single export operation.

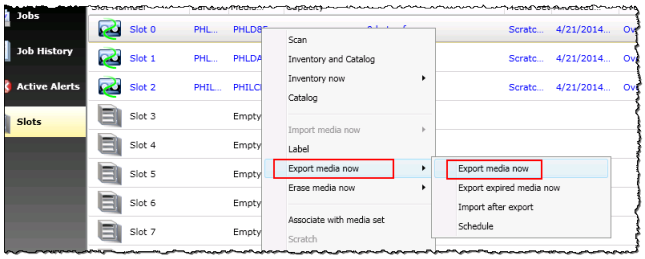


2. In the **Media Request** pop-up window, click **Details**, and then click **Respond OK** in the **Alert: Media Intervention** window.

   In the AWS Storage Gateway console, you can verify the status of the tape you are archiving. It might take some time to finish uploading data to AWS. During this time, the exported tape will be listed in the gateway's VTL with the status IN TRANSIT TO VTS. When the upload is completed and the archiving process begins, the status changes to ARCHIVING. When data archiving has completed, the exported tape will no longer be listed in the VTL.
3. Click your gateway, and then click **VTL Tape Cartridges** and verify that the virtual tape is no longer listed in your gateway.
4. On the Navigation pane of the AWS Storage Gateway console, click **Virtual Tape Shelf (VTS)**. Verify that your archived tape is listed in the VTS and that the status is ARCHIVED.

# Restoring Data from a Tape Archived in Backup Exec 2012

Restoring your archived data is a two-step process.

### To restore data from an archived tape

1. Retrieve the archived tape from your VTS to a gateway-VTL. For instructions, see Retrieving the Archived Tape from the VTS Back to Your Gateway-VTL (p. 260).
2. Use Backup Exec 2012 to restore the data. This process is the same as restoring data from physical tapes. For instructions, go to the Backup Exec documentation on the Symantec website.

# Disabling a Tape Drive in Backup Exec 2012

AWS Storage Gateway-VTL provides 10 tape drives, but you might decide to use fewer tape drives. In that case, you disable the tape drives you will not use.

1. Open Backup Exec 2012, and click the **Storage** tab.
2. In the **Robotic library** tree, right-click the tape drive you want to disable, and then click **Disable**.

# Document History for AWS Storage Gateway

The following table describes the important changes to the documentation since the last release of the *AWS Storage Gateway User Guide*.

- **API version**: 2013-06-30
- **Latest documentation update**: June 22, 2015

| Change | Description | Date Changed |
| --- | --- | --- |
| Support for Dell NetVault Backup 10.0 | AWS Storage Gateway-VTL now supports Dell NetVault Backup 10.0. You can now use Dell NetVault Backup 10.0 to back up your data to Amazon S3 and archive directly to Amazon Glacier. For more information, see Testing Your Setup by Using Dell NetVault Backup (p. 273). | In this release. |

| Change | Description | Date Changed |
| --- | --- | --- |
| Support for 16 TiB storage volumes for Gateway-stored setups<br><br>Support for system resource checks on the AWS Storage Gateway local console<br><br>Support for the Red Hat Enterprise Linux 6 iSCSI initiator | AWS Storage Gateway now supports 16 TiB storage volumes for Gateway-stored setups. You can now create 12 16 TiB storage volumes for a maximum of 192 TiB of storage. For more information, see Gateway-Stored Volume Architecture (p. 4).<br><br>You can now determine whether your system resources (virtual CPU cores, root volume size, and RAM) are sufficient for your gateway to function properly. For more information, see Viewing Your Gateway System Resource Errors (p. 150) or Viewing Your Gateway-VTL System Resource Errors (p. 334).<br><br>AWS Storage Gateway now supports the Red Hat Enterprise Linux 6 iSCSI initiator. For more information, see Requirements (p. 11).<br><br>This release includes the following AWS Storage Gateway improvements and updates:<br><br>• From the AWS Storage Gateway console, you can now see the date and time the last successful software update was applied to your gateway. For more information, see Managing Gateway Updates Using the AWS Storage Gateway Console (p. 129) or Managing Gateway-VTL Updates Using the AWS Storage Gateway Console (p. 312)<br>• AWS Storage Gateway now provides an API you can use to list iSCSI initiators connected to your storage volumes. For more information, see ListVolumeInitiators in the API reference. | June 03, 2015 |
| Support for Microsoft Hyper-V hypervisor versions 2012 and 2012 R2 | AWS Storage Gateway now supports Microsoft Hyper-V hypervisor versions 2012 and 2012 R2. This is in addition to support for Microsoft Hyper-V hypervisor version 2008 R2. For more information, see Supported Hypervisors and Host Requirements (p. 11). | April 30, 2015 |
| Support for Symantec Backup Exec 15 | AWS Storage Gateway-VTL now supports Symantec Backup Exec 15. You can now use Symantec Backup Exec 15 to back up your data to Amazon S3 and archive directly to Amazon Glacier. For more information, see Testing Your Setup by Using Symantec Backup Exec (p. 263). | April 06, 2015 |
| CHAP authentication support for storage volumes | AWS Storage Gateway now supports configuring CHAP authentication for storage volumes. For more information, see Step 3: Create Volumes (p. 74). | April 02, 2015 |

| Change | Description | Date Changed |
| --- | --- | --- |
| Support for VMware ESXi Hypervisor version 5.1 and 5.5 | AWS Storage Gateway now supports VMware ESXi Hypervisor versions 5.1 and 5.5. This is in addition to support for VMware ESXi Hypervisor versions 4.1 and 5.0. For more information, see Supported Hypervisors and Host Requirements (p. 11). | March 30, 2015 |
| Support for Windows CHKDSK utility | AWS Storage Gateway now supports the Windows CHKDSK utility. You can use this utility to verify the integrity of your volumes and fix errors on the volumes. For more information, see Troubleshooting Storage Volume Issues (p. 233). | March 04, 2015 |
| Integration with AWS CloudTrail to capture API calls | AWS Storage Gateway is now integrated with AWS CloudTrail. AWS CloudTrail captures API calls made by or on behalf of AWS Storage Gateway in your AWS account and delivers the log files to an Amazon S3 bucket that you specify. For more information, see Logging Volume Gateway API Calls by Using AWS CloudTrail (p. 200).<br><br>This release includes the following AWS Storage Gateway improvement and update:<br><br>• Virtual tapes that have dirty data in cache storage (that is, that contain content that has not been uploaded to AWS) are now recovered when a gateway cached drive changes. For more information, see Recovering a Virtual Tape (p. 282). | December 16, 2014 |
| Support for additional backup software and medium changer | AWS Storage Gateway-VTL now supports the following backup software:<br><br>• Symantec Backup Exec 2014<br>• Microsoft System Center 2012 R2 Data Protection Manager<br>• Veeam Backup & Replication V7<br>• Veeam Backup & Replication V8<br><br>You can now use these four backup software products with the Storage Gateway virtual tape library (VTL) to back up to Amazon S3 and archive directly to Amazon Glacier. For more information, see Step 5: Test the Gateway-VTL Setup (p. 249).<br><br>AWS Storage Gateway now provides an additional medium changer that works with the new backup software.<br><br>This release includes miscellaneous AWS Storage Gateway-VTL improvements and updates. | November 3, 2014 |

| Change | Description | Date Changed |
| --- | --- | --- |
| EU (Frankfurt) region | AWS Storage Gateway is now available in the EU (Frankfurt) region. For detailed information, see Available AWS Regions (p. 7). | October 23, 2014 |
| Content restructure | Created a Getting Started section that is common to all gateway solutions. This section provides instructions for you to download, deploy, and activate a gateway. After you deploy and activate a gateway, you can proceed to further instructions specific to gateway-stored, gateway-cached, and gateway-VTL setups. For more information, see Getting Started with AWS Storage Gateway (p. 9). | May 19, 2014 |
| Support for Symantec Backup Exec 2012 | AWS Storage Gateway-VTL now supports Symantec Backup Exec 2012. You can now use Symantec Backup Exec 2012 to back up your data to Amazon S3 and archive directly to Amazon Glacier. For more information, see Appendix I: Testing Your Setup by Using Symantec Backup Exec 2012 (p. 426). | April 28, 2014 |
| Support for Windows Server Failover Clustering<br><br>Support for VMware ESX initiator<br><br>Support for performing configuration tasks on AWS Storage Gateway local console | • Gateway-cached and gateway-stored volumes can now be used as storage for applications configured using Windows Server Failover Clustering (WSFC). This enables coordinated iSCSI access to AWS Storage Gateway volumes by applications clustered using WSFC.<br><br>• AWS Storage Gateway now enables you to manage storage connectivity directly through your ESX host. This provides an alternative to using initiators resident in the guest OS of your VMs.<br><br>• AWS Storage Gateway now provides support for performing configuration tasks in the AWS Storage Gateway local console. For information about performing configuration tasks on gateways deployed on-premises, see Maintenance Tasks When Running a Gateway On-Premises (p. 137) or Maintenance Tasks When Running a Gateway-VTL On-Premises (p. 321). For information about performing configuration tasks on gateways deployed on an EC2 instance, see Performing Maintenance Tasks on the Amazon EC2 Gateway Local Console (p. 162) or Performing Maintenance Tasks on the Local Console for an Amazon EC2 Gateway-VTL (p. 345). | January 31, 2014 |

| Change | Description | Date Changed |
| --- | --- | --- |
| Support for virtual tape library (VTL) and introduction of API version 2013-06-30 | AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to integrate your on-premises IT environment with the AWS storage infrastructure. In addition to volume gateways (gateway-cached and gateway-stored), AWS Storage Gateway now supports gateway–virtual tape library (VTL). You can configure gateway-VTL with up to 10 virtual tape drives per gateway. Each virtual tape drive responds to the SCSI command set, so your existing on-premises backup applications will work without modification. For more information, see the following topics in the *AWS Storage Gateway User Guide*.<br>• For an architectural overview, see Gateway–Virtual Tape Library (Gateway-VTL) Architecture (p. 5).<br>• To get started with gateway-VTL, see Getting Started with AWS Storage Gateway (p. 9). | November 5, 2013 |
| Support for Microsoft Hyper-V | AWS Storage Gateway now provides the ability to deploy an on-premises gateway on the Microsoft Hyper-V virtualization platform. Gateways deployed on Microsoft Hyper-V have all the same functionality and features as the existing on-premises storage gateway. To get started deploying a gateway with Microsoft Hyper-V, see Hyper-V Host (p. 35). | April 10, 2013 |
| Support for deploying a gateway on Amazon EC2 | AWS Storage Gateway now provides the ability to deploy a gateway in Amazon Elastic Compute Cloud (Amazon EC2). You can launch a gateway instance in Amazon EC2 using the AWS Storage Gateway AMI available in AWS Marketplace. To get started deploying a gateway using the AWS Storage Gateway AMI, go to Amazon EC2 Gateway (p. 57). | January 15, 2013 |

| Change | Description | Date Changed |
| --- | --- | --- |
| Support for gateway-cached volumes and introduction of API Version 2012-06-30 | In this release, AWS Storage Gateway introduces support for gateway-cached volumes. Gateway-cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their active data. You can create storage volumes up to 32 TiB in size and mount them as iSCSI devices from your on-premises application servers. Data written to your gateway-cached volumes is stored in Amazon Simple Storage Service (Amazon S3), with only a cache of recently written and recently read data stored locally on your on-premises storage hardware. Gateway-cached volumes allow you to utilize Amazon S3 for data where higher retrieval latencies are acceptable, such as for older, infrequently accessed data, while maintaining storage on-premises for data where low-latency access is required.<br><br>In this release, AWS Storage Gateway also introduces a new API version that, in addition to supporting the current operations, provides new operations to support gateway-cached volumes.<br><br>For more information on the two AWS Storage Gateway solutions, see How AWS Storage Gateway Works (Architecture) (p. 2).<br><br>You can also try a test setup. For instructions, go to Getting Started with AWS Storage Gateway (p. 9). | October 29, 2012 |
| API and IAM Support | In this release, AWS Storage Gateway introduces API support as well as support for AWS Identity and Access Management(IAM).<br><br>• **API support—**You can now programmatically configure and manage your AWS Storage Gateway resources. For more information about the APIs, see API Reference for AWS Storage Gateway (p. 379) in the *AWS Storage Gateway User Guide*.<br>• **IAM Support** – AWS Identity and Access Management (IAM) enables you create users and manage user access to your AWS Storage Gateway resources by means of IAM policies. For examples of IAM policies, go to Access Control Using AWS Identity and Access Management (IAM) (p. 373). For more information about IAM, go to AWS Identity and Access Management (IAM) detail page. | May 9, 2012 |
| Static IP Support | You can now specify a static IP for your local gateway. For more information, see Configuring Your Gateway to Use a Static IP Address (p. 142). | March 5, 2012 |
| New Guide | This is the first release of *AWS Storage Gateway User Guide*. | January 24, 2012 |